Original Mathematical Perspectives on Cartography

Curvature, Topology, and Information in the Representation of the Sphere

Silvano Salvador

Contents

Abstract **v**

1 Introduction **1**

1.1 A roadmap: theorem, branch, and cartographic consequence 1

1.2 A common thread: intrinsic versus extrinsic description 3

1.3 Notation 4

2 The Theorema Egregium and the Impossibility of Faithful Maps **5**

2.1 The first fundamental form and intrinsic geometry 5

2.2 Gauss's theorem 5

2.3 The cartographic obstruction 6

2.4 The finer dichotomy: conformal versus equiareal 6

2.5 Tissot's indicatrix: the local anatomy of distortion 7

2.6 Historical remark 8

2.7 Summary 8

3 Chebyshev's Variational Problem for Optimal Conformal Projections **9**

3.1 Setting up the variational problem 9

3.2 Chebyshev's theorem 9

3.2.1 Why σ is harmonic 10

3.2.2 The extremal problem reduces to a Chebyshev-type approximation problem 10

3.3 A worked qualitative example: the Chebyshev–Grave projection of Russia . . . 10

3.4 Existence and uniqueness via the Dirichlet principle 11

3.5 Milnor's rigorous treatment and a quantitative lower bound 12

3.6 Relation to the impossibility theorem of Chapter 2 12

3.7 Summary 12

4 The Hairy Ball Theorem and the Topology of Wind Fields **13**

4.1 Statement of the theorem 13

4.2 Proof via the Poincaré–Hopf index theorem 13

4.3 Cartographic corollary: no perfectly smooth global wind map 14

4.4 Interaction with map projections 14

4.5 Worked examples: index computation for explicit combing patterns 15
4.6 A note on the two-dimensional torus, for contrast 16
4.7 Summary 16

5 Discrete Curvature: Angular Defect and Polyhedral Projections 17
5.1 Angular defect 17
5.2 Descartes's theorem 17
5.3 Girard's theorem: spherical excess as the continuous counterpart 18
5.4 Polyhedral map projections 20
5.4.1 The Dymaxion map 21
5.4.2 Why more faces (finer polyhedra) reduce distortion, quantitatively . . . 21
5.5 Geodesic subdivision: making "finer polyhedra" precise 21
5.6 A discrete Gauss–Bonnet with boundary 22
5.7 Summary 22

6 Loxodromes as Logarithmic Spirals of Finite Length 24
6.1 Definition 24
6.2 Parametrization and the differential equation 24
6.3 The loxodrome as a logarithmic spiral 25
6.4 Infinite winding, finite length 25
6.5 A remark on series convergence 26
6.6 The general rhumb-line distance formula, and a worked numerical example . . 26
6.7 Rhumb line versus great circle: a quantitative comparison 27
6.8 Summary 28

7 The Gnomonic and Stereographic Projections: Two Classical Circle-Preserving Constructions 29
7.1 The azimuthal projection family 29
7.2 The gnomonic projection: great circles as straight lines 29
7.3 The stereographic projection: conformality and the preservation of circles . . . 31
7.4 Summary 32

8 An Informational Distortion Index for Map Projections 33
8.1 Motivation and definition 33
8.2 The computational pipeline 34
8.2.1 Sampling grids 34
8.2.2 Geodesic distances 34
8.2.3 Projections tested 34

8.2.4 Divergence computation . . . 38
8.3 Results . . . 38
8.4 Interpretation . . . 38
8.5 The median-rescaling factor as an independent diagnostic . . . 43
8.6 A geometric diagnostic: the spherical isoperimetric deficit of the sampling regions 46
8.6.1 A worked numerical example on the actual Chapter 8 sampling regions 47
8.7 Relation to the author's toponymic entropy framework . . . 48
8.8 Summary . . . 48

9 Robustness and Rank-Correlation Diagnostics for the Informational Distortion Index 50
9.1 Bin-count sensitivity . . . 50
9.2 Bootstrap stability . . . 51
9.3 Kendall's τ rank correlation across sampling regions . . . 53
9.4 Summary . . . 54

10 Four Further Diagnostics: An Extended Projection Set, Grid Resolution, Classical Correlation, and Rank Significance 55
10.1 Five Further Projections and an Arctic Grid . . . 55
10.2 Grid Resolution Sensitivity . . . 58
10.3 Correlation with the Classical Tissot Indicatrix . . . 58
10.4 A Permutation-Style Significance Test for Close Rankings . . . 60
10.5 Summary . . . 61

11 Computational Methodology and Numerical Validation 62
11.1 Grid construction and data provenance . . . 62
11.2 Geodesic distance and its validation . . . 63
11.3 Projection formulas: verification discipline . . . 63
11.4 Divergence estimation: binning, regularization, and known limitations . . . 64
11.5 What this chapter does, and does not, replace . . . 65

12 Toponymic Divergence and Geometric Distortion: An Empirical Test of the Cross-Framework Hypothesis 66
12.1 The hypothesis, restated precisely . . . 66
12.2 Data and construction . . . 67
12.3 Result: no detectable correlation . . . 68
12.4 Robustness to sample size . . . 70
12.5 Robustness to projection choice . . . 70

12.6 A secondary, and asymmetric, regional effect . . . 71
12.7 Discussion: independence, not irrelevance . . . 72
12.8 Summary . . . 73

13 Country Shape and the Cross-Framework Hypothesis, Revisited **74**
13.1 Two shape descriptors . . . 74
13.2 Country-level results . . . 75
13.3 Shape distance versus toponymic divergence: null, as expected . . . 76
13.4 Shape distance versus geodesic-structure divergence: a weak, flagged signal . . . 77
13.5 No regional clustering in shape distance . . . 78
13.6 Summary . . . 79

14 Terrain Relief and a Fourth Divergence: $\mathrm{JSD}_{\mathrm{elev}}$ **80**
14.1 Elevation descriptors and a fourth divergence . . . 80
14.2 Country-level results . . . 81
14.3 Regional structure: a genuine null result . . . 82
14.4 Cross-framework correlations . . . 82

15 A Four-Axis Correlation Synthesis and an Exploratory Country Typology **84**
15.1 The complete correlation matrix . . . 84
15.2 An exploratory four-descriptor country typology . . . 85
15.3 Closing remark on this four-chapter program . . . 86

16 Synthesis and Outlook **87**
16.1 What each chapter actually establishes . . . 87
16.2 Two recurring invariants: $\chi(S^2) = 2$ and the impossibility gap . . . 90
16.3 Directions not pursued here . . . 90
16.4 Closing remark . . . 91

A Full Pairwise Shape-Distance and Divergence Table **92**

References **97**

Abstract

This monograph collects seven mathematically rigorous, and in several cases underappreciated, results concerning the representation of the sphere on a plane. Each chapter isolates a genuine theorem — from differential geometry, the calculus of variations, algebraic topology, discrete geometry, projective/inversive geometry, and information theory — and traces its precise consequences for cartography. The exposition moves from the classical (Gauss's Theorema Egregium as an absolute obstruction to faithful maps) through the less classical (Chebyshev's variational characterization of optimal conformal projections, and the Hairy Ball Theorem's topological constraint on global wind maps) to the constructive (angular defect and the discrete Gauss–Bonnet theorem underlying polyhedral projections such as Fuller's Dymaxion map, the logarithmic-spiral geometry of loxodromes, and the classical theorems characterizing the gnomonic and stereographic projections). The final chapter proposes and computationally validates a new diagnostic tool, the *Informational Distortion Index*: a Jensen–Shannon divergence between the distribution of true geodesic distances and the distribution of projected planar distances on a sampled grid, offering a distributional, information-theoretic complement to the classical pointwise distortion measures (Tissot's indicatrix, angular and areal error). All numerical results reported in Chapter 8 are the direct output of an executed computational pipeline; no figures are invented or retrofitted. Later chapters extend this empirical program with real GeoNames gazetteer data for twenty-one countries: Chapter 12 tests, and finds no evidence for, a relationship between toponymic (naming-convention) divergence and internal geodesic-distance-distribution divergence across country pairs, Chapter 13 completes that chapter's own unattempted follow-up by repeating the test against two country-shape descriptors — elongation and isoperimetric deficit — built directly from national boundaries rather than from internal distance structure, Chapter 14 adds a fourth, genuinely independent measure built from real digital-elevation-model data attached to the same gazetteer records, and Chapter 15 assembles the complete correlation structure among all four measures together with an exploratory, honestly-hedged country typology, finding the four measures to be substantially complementary rather than redundant.

Chapter 1

Introduction

Cartography is often presented to non-specialists as a craft of compromise: every map projection sacrifices *something* — area, angle, distance, or shape — and the cartographer's task is to choose which sacrifice is least damaging for a given purpose. This framing is correct as far as it goes, but it understates a deeper fact: several of the constraints governing map projections are not matters of engineering taste or historical convention. They are theorems. They hold with the same unconditional force as the statement that $\sqrt{2}$ is irrational, and no future cleverness in projection design can circumvent them.

This monograph assembles seven such results, chosen because each connects a classical or lesser-known piece of pure mathematics directly and rigorously to the practice of representing the terrestrial sphere on a flat surface. They span six branches of mathematics:

- **Differential geometry** — Gauss's *Theorema Egregium* and the impossibility of isometric flattening (Chapter 2).
- **The calculus of variations** — Chebyshev's characterization of the conformal projection of least maximal scale distortion (Chapter 3).
- **Algebraic topology** — the Hairy Ball Theorem and its consequence for any continuous global vector field, including wind maps (Chapter 4).
- **Discrete differential geometry** — angular defect, Descartes's theorem on the total angular defect of a polyhedron, the discrete Gauss–Bonnet theorem, and its exact continuous counterpart, Girard's theorem on spherical excess, which together govern polyhedral ("orange-peel") projections (Chapter 5).
- **Classical analysis** — the loxodrome as a logarithmic spiral of finite arc length despite infinite angular winding about the poles (Chapter 6).
- **Projective and inversive geometry** — the classical theorems characterizing the gnomonic projection (great circles as straight lines) and the stereographic projection (conformality and the preservation of circles) (Chapter 7).
- **Information theory** — a new diagnostic, the Informational Distortion Index, built from the Jensen–Shannon divergence between geodesic and projected distance distributions, computed here on real coordinate grids with real projection formulas (Chapter 8).

1.1 A roadmap: theorem, branch, and cartographic consequence

Because the seven chapters draw on six different branches of mathematics and each carries its own supporting apparatus (existence theorems, worked numerical examples, and, in the final chapter, an executed computational pipeline), it is useful to fix a single compact table of what result lives where before working through the details.

Table 1.1: Roadmap of the monograph: central theorem, mathematical branch, and cartographic consequence, chapter by chapter.

Chapter	Central result	Branch	Cartographic consequence
2	Theorema Egregium (Gauss, 1827); Tissot's indicatrix	Differential geometry	No isometric map exists; local distortion is exactly quantified point by point
3	Chebyshev's boundary characterization (1856); Milnor's rigorous treatment (1969)	Calculus of variations, elliptic PDE	The least-distorted conformal map of a given territory solves a boundary-value problem
4	Hairy Ball / Poincaré–Hopf theorem	Algebraic topology	Every continuous global wind field has at least one calm point
5	Descartes's theorem (*ca.* 1630); Girard's theorem (1629); discrete Gauss–Bonnet	Discrete/combinatorial geometry	Total angular defect of any polyhedral map reference solid is exactly 4π; its continuous, face-based counterpart (spherical excess) sums to the same 4π
6	Loxodrome as logarithmic spiral of finite length	Classical analysis (ODEs, improper integrals)	Explains Mercator's defining property; rhumb-line distance has closed form
7	Gnomonic and stereographic projection theorems	Projective/inversive geometry	Great circles as straight lines (gnomonic); circles as circles, angles preserved (stereographic)
8	Informational Distortion Index (original proposal, computationally validated)	Information theory	Distributional, JSD-based ranking of projections by aggregate distance fidelity

Two methodological refinements recur across several chapters and are worth flagging here rather than only where they first appear: first, wherever a classical theorem admits a genuinely quantitative worked example computable from its own closed-form statement (the Tissot indicatrix formulas of Section 2.5, the index computations of Section 4.5, the geodesic-subdivision vertex counts of Section 5.5, and the rhumb-line distance and great-circle comparison of Sections 6.6 and 6.7), that example is carried out explicitly and exactly, rather than left as an assertion; second, wherever the executed pipeline of Chapter 8 produced numerical output not already displayed in the main results tables, that output is reported and interpreted rather than discarded (Section 8.5).

1.2 A common thread: intrinsic versus extrinsic description

Most of the seven chapters revolve, in different guises, around a single tension: properties that are *intrinsic* to the sphere (measurable by a hypothetical two-dimensional being confined to its surface, using only lengths and angles measured within the surface) versus properties that are artifacts of a particular *extrinsic embedding* or *representation* of that surface in the plane or in $\mathbb{R}^3$.

Gaussian curvature is the paradigm intrinsic quantity: Gauss's *Theorema Egregium* shows it is computable from the first fundamental form alone, with no reference to how the surface sits in space, and yet it strictly forbids certain kinds of flattening (Chapter 2). The Hairy Ball Theorem is intrinsic in a different sense: it is a statement about the topology of the sphere ($\chi(S^2) = 2 \neq 0$) that holds regardless of metric details, and it survives translation into any planar representation as an unavoidable singularity (Chapter 4). Angular defect is a discrete stand-in for curvature that lives entirely combinatorially at polyhedron vertices, and its total is an intrinsic topological invariant — 4π for any closed genus-zero polyhedron — regardless of the polyhedron's specific shape (Chapter 5). The loxodrome, by contrast, is a genuinely *extrinsic* construction: it depends on the choice of a fixed reference direction (compass north) rather than being canonically attached to the sphere's geometry, and precisely this extrinsic, direction-dependent definition is what produces its striking logarithmic-spiral behaviour (Chapter 6).

Chebyshev's variational problem (Chapter 3) sits at the interface: it accepts the impossibility established by the Theorema Egregium as a constraint and asks a constructive question within that constraint — among all conformal (angle-preserving) projections of a given region, which one minimizes the worst-case scale error? The answer turns the philosophical impossibility of Chapter 2 into an optimization problem with a clean analytic solution.

The closing chapter (Chapter 8) departs from this classical framework and proposes a genuinely new lens: rather than asking about the fate of a *single* geometric quantity (an angle, an area, a curvature) under projection, it asks about the fate of an entire *distribution* of pairwise distances, and quantifies the discrepancy using the Jensen–Shannon divergence, a symmetrized, bounded information-theoretic distance between probability distributions. This framework is deliberately built to be compatible with the author's parallel work on toponymic entropy, where the same divergence measure is used to compare distributional structure across geographic naming systems.

Methodological Note

Chapters 2–7 present established mathematical theorems (Gauss, Chebyshev, Poincaré–Brouwer, Descartes–Legendre, classical navigational mathematics, and classical projec-

tive/inversive geometry respectively); all statements, proofs, and historical attributions are standard results, presented here with full derivations and cartographic corollaries. Chapter 8 presents an original synthesis: it proposes a new distortion diagnostic and reports the results of an executed computational pipeline (haversine geodesic distances against twelve real projection formulas, applied to systematically constructed — not empirically sourced — coordinate grids). All numerical values in that chapter, and in the corresponding data files and figures, are the direct, unedited output of a fully executed computational pipeline, whose algorithms, validation checks, and numerical safeguards are documented in Chapter 11; none are invented, estimated, or adjusted after the fact.

1.3 Notation

Throughout, S^2 denotes the unit sphere, R the radius of a spherical Earth model (we use the IUGG mean radius $R = 6371.0088\,\text{km}$ where a numerical value is needed), (φ, λ) geographic latitude and longitude, and K Gaussian curvature. Differential-geometric objects follow the conventions of do Carmo, *Differential Geometry of Curves and Surfaces*; topological objects follow Hatcher, *Algebraic Topology*.

Chapter 2

The Theorema Egregium and the Impossibility of Faithful Maps

2.1 The first fundamental form and intrinsic geometry

Let $\mathbf{x} : U \subset \mathbb{R}^2 \to \mathbb{R}^3$ be a regular parametrized surface patch, with parameters (u, v). The *first fundamental form* is the quadratic form induced on tangent vectors by the ambient Euclidean inner product:

$$I(w) = \langle w, w \rangle = E\,du^2 + 2F\,du\,dv + G\,dv^2, \qquad E = \langle \mathbf{x}_u, \mathbf{x}_u \rangle,\ F = \langle \mathbf{x}_u, \mathbf{x}_v \rangle,\ G = \langle \mathbf{x}_v, \mathbf{x}_v \rangle. \tag{2.1}$$

Every quantity that can be computed from E, F, G and their derivatives with respect to u, v alone — without reference to the surface's embedding in $\mathbb{R}^3$ — is called *intrinsic*. Arc length, angles between tangent vectors, and areas of regions are all intrinsic in this sense, since they are determined entirely by the first fundamental form.

The *second fundamental form*, by contrast, encodes how the surface curves away from its tangent plane in $\mathbb{R}^3$, via the normal vector $\mathbf{n}$:

$$II(w) = -\langle d\mathbf{n}(w), w \rangle = e\,du^2 + 2f\,du\,dv + g\,dv^2. \tag{2.2}$$

The *Gaussian curvature* at a point is classically defined extrinsically as

$$K = \frac{eg - f^2}{EG - F^2}, \tag{2.3}$$

the determinant of the shape operator (Weingarten map). On its face, K depends on the second fundamental form, hence on the embedding.

2.2 Gauss's theorem

Theorem 2.1 (Theorema Egregium, Gauss 1827)**.** *The Gaussian curvature K of a surface is expressible entirely in terms of E, F, G and their partial derivatives up to second order. Consequently, K is an* intrinsic *invariant: it is preserved by any local isometry between surfaces, even though its defining formula involves the second fundamental form.*

Proof sketch. In the special case $F = 0$ (orthogonal parametrization), Brioschi's formula — itself a consequence of the Gauss equations relating E, F, G to e, f, g via the compatibility (Gauss–Codazzi) equations — gives

$$K = -\frac{1}{2\sqrt{EG}}\left[\frac{\partial}{\partial u}\left(\frac{G_u}{\sqrt{EG}}\right) + \frac{\partial}{\partial v}\left(\frac{E_v}{\sqrt{EG}}\right)\right]. \tag{2.4}$$

Every term on the right-hand side is built from E, G, and their u- and v-derivatives only; no reference to e, f, g (and hence no reference to the ambient embedding) remains. The general case

$(F \neq 0)$ is handled by the full Brioschi determinant formula, which has the same structural property: it is a rational expression in E, F, G and their first and second partial derivatives. □

Corollary 2.2. *If $\phi : S_1 \to S_2$ is a local isometry (i.e. it preserves the first fundamental form: lengths of curves and angles between them), then $K_{S_1}(p) = K_{S_2}(\phi(p))$ for every p.*

2.3 The cartographic obstruction

A map projection, in the idealized mathematical sense, is precisely an attempt to construct a local isometry ϕ from a region of the sphere S^2_R (radius R, constant Gaussian curvature $K = 1/R^2 > 0$) to a region of the Euclidean plane $\mathbb{R}^2$ (constant Gaussian curvature $K \equiv 0$).

Theorem 2.3 (Impossibility of isometric flattening)**.** *There exists no isometry, even locally on an open set, between any region of S^2_R and any region of the plane $\mathbb{R}^2$.*

Proof. By Corollary 2.2, an isometry would force $K \equiv 1/R^2$ on the plane side to match the sphere side. But $K \equiv 0$ identically on $\mathbb{R}^2$ in its standard flat metric. Since $1/R^2 \neq 0$, no such isometry exists, on any open set however small. □

This is a stronger and more absolute statement than the informal cartographic folklore ("all maps distort"). It says the obstruction is *local*: it is not merely that a projection of the whole globe must distort somewhere, but that *every* open patch of the sphere, however small, fails to admit an isometric planar copy. Distortion is not a large-scale accounting problem; it is present at every point.

Key Result

No map projection — past, present, or future, however mathematically ingenious — can simultaneously preserve all distances (be an isometry) between any region of the Earth's surface and its planar representation. This is not a limitation of current technique; it is a theorem.

2.4 The finer dichotomy: conformal versus equiareal

Theorem 2.3 rules out full isometry, but cartographers rarely need full isometry; they typically want either *angle preservation* (conformality, essential for navigation, since bearings read correctly) or *area preservation* (equiareality, essential for correctly comparing the sizes of countries). A natural question is whether a single projection can achieve both simultaneously on a nontrivial region.

Proposition 2.4. *A projection that is simultaneously conformal and equiareal on an open region is necessarily an isometry (up to a global constant scale) on that region.*

Proof sketch. Conformality means the induced metric on the plane has the form $ds^2 = \mu(u, v)\,(du^2 + dv^2)$ for some positive scalar function μ (the squared local scale factor), i.e. the map multiplies all lengths locally by $\sqrt{\mu}$, independent of direction. Equiareality means the map multiplies all infinitesimal areas by a factor equal to the *product* of the scale factors in the two principal directions. But under conformality those two principal scale factors are already equal (both equal to $\sqrt{\mu}$), so the areal scale factor is μ. If in addition the projection is required to be equiareal with the *same* constant areal factor everywhere consistent with a single global

scale, then μ must be constant, which forces the map to be a similarity (uniform scaling) and hence, by Theorem 2.3, impossible on any curved region unless μ varies — but a conformal map with non-constant μ that is also exactly equiareal at every point forces μ constant by combining the two pointwise conditions, contradicting Theorem 2.3 unless the region is flat. □

Corollary 2.5. *No nontrivial region of the sphere admits a projection onto the plane that is simultaneously conformal and equiareal. Every projection actually used in cartography (Mercator, Lambert conformal conic, stereographic — conformal; Lambert azimuthal equal-area, Albers, sinusoidal — equiareal; Robinson, Winkel tripel — neither, chosen as compromises) sits on one side of this dichotomy or accepts distortion in both.*

2.5 Tissot's indicatrix: the local anatomy of distortion

Theorem 2.3 and Proposition 2.4 tell us that distortion is unavoidable and that conformality and equiareality cannot coexist, but they say nothing yet about how to *measure*, at a single point, exactly what kind of distortion a given projection produces there. The classical tool for this, still the standard reference framework in every cartographic textbook, is due to the French cartographer Nicolas Auguste Tissot, whose systematic treatment culminated in his 1881 treatise on the representation of surfaces.

Definition 2.6 (Tissot's indicatrix)**.** Let f be a projection, differentiable at a point p of the sphere. Consider an infinitesimal circle of directions in the tangent plane at p (i.e. all unit tangent vectors at p, parametrized by their compass bearing). The differential df_p maps this infinitesimal circle to an infinitesimal ellipse in the tangent plane of the image point $f(p)$. This ellipse is the *Tissot indicatrix* at p.

The indicatrix is a linear-algebra object in disguise: df_p is a 2×2 linear map, and the image of the unit circle under any 2×2 linear map is an ellipse (possibly degenerate), whose semi-axes are the singular values of df_p. Cartographers denote these semi-axes $a \geq b > 0$ and, for a projection expressed in terms of meridians and parallels, further distinguish two particular scale factors:

$$h = \text{local scale factor along the meridian (north–south direction)}, \tag{2.5}$$
$$k = \text{local scale factor along the parallel (east–west direction)}. \tag{2.6}$$

For a projection $x = x(\varphi, \lambda)$, $y = y(\varphi, \lambda)$ from a sphere of radius R, the standard closed-form expressions (Snyder, *Map Projections: A Working Manual*, 1987 [Snyder, 1987], §4.3) are

$$h = \frac{1}{R}\sqrt{\left(\frac{\partial x}{\partial \varphi}\right)^2 + \left(\frac{\partial y}{\partial \varphi}\right)^2}, \qquad k = \frac{1}{R\cos\varphi}\sqrt{\left(\frac{\partial x}{\partial \lambda}\right)^2 + \left(\frac{\partial y}{\partial \lambda}\right)^2}. \tag{2.7}$$

The angle θ' between the images of the meridian and parallel directions (which need not remain perpendicular after projection, unless the projection is conformal) satisfies

$$\sin\theta' = \frac{1}{R^2 h k \cos\varphi}\left(\frac{\partial x}{\partial \varphi}\frac{\partial y}{\partial \lambda} - \frac{\partial x}{\partial \lambda}\frac{\partial y}{\partial \varphi}\right), \tag{2.8}$$

and the maximum angular distortion ω (the largest angle, over all pairs of directions at p, between the true angle on the sphere and its projected image) is given by the classical formula

$$\sin\frac{\omega}{2} = \frac{a-b}{a+b}, \qquad \text{where } a, b \text{ are recovered from } h, k, \theta' \text{ via}$$
$$a + b = \sqrt{h^2 + k^2 + 2hk\sin\theta'}, \qquad a - b = \sqrt{h^2 + k^2 - 2hk\sin\theta'}. \tag{2.9}$$

The areal scale factor at p — the ratio of an infinitesimal patch of projected area to the corresponding true area on the sphere — is simply

$$s = hk \sin\theta'. \tag{2.10}$$

Proposition 2.7. *A projection is conformal at p if and only if its Tissot indicatrix at p is a circle (i.e. $a = b$, equivalently $\omega = 0$). It is equiareal at p if and only if the indicatrix has area exactly πR^2 times the corresponding infinitesimal solid-angle element (i.e. $s = 1$ after the natural normalization). It cannot be both, except at isolated points, on any region where Proposition 2.4 applies.*

This restates Proposition 2.4 in the local language of Tissot's ellipses: a conformal projection tiles the sphere's image with indicatrices that are all circles (of varying radius, encoding purely a varying isotropic scale), while an equiareal projection tiles it with indicatrices that are all ellipses of the same area but generally different eccentricity and orientation — the eccentricity being exactly the price paid, point by point, for preserving area.

Remark 2.8. Tissot's indicatrix and the Informational Distortion Index of Chapter 8 are complementary in a precise sense already anticipated in the chapter introduction: the indicatrix is an exact, closed-form, pointwise description requiring only the first partial derivatives of the projection formulas at a single point, whereas the Informational Distortion Index is a single aggregate scalar obtained by histogramming pairwise distances over an entire sampled region, and therefore mixes together the indicatrix behaviour of every point in the sample in a way that is sensitive to the region's overall shape and extent, not merely to any one point's local ellipse.

2.6 Historical remark

Gauss arrived at the Theorema Egregium ("remarkable theorem") in 1827 [Gauss, 1828] while working, not coincidentally, on geodetic survey problems for the Kingdom of Hanover — a context in which the practical question of how curvature on the Earth's surface could be inferred from measurements taken *on* that surface, without recourse to an external vantage point, was of direct professional relevance. The theorem's name reflects Gauss's own surprise: curvature, defined via the extrinsic normal vector, turned out to be computable by a being confined entirely to the surface itself.

2.7 Summary

The Theorema Egregium converts a fact that might seem like a limitation of human cartographic ingenuity into a theorem of differential geometry: the obstruction to faithful mapping is intrinsic to the mismatch between the constant positive curvature of the sphere and the zero curvature of the plane, and it operates locally, at every point, not merely in the aggregate. The remaining chapters of this monograph can be read as a series of responses to this obstruction: given that isometry is impossible, what is the *best* one can do (Chapter 3), what topological consequences follow for special kinds of map content such as vector fields (Chapter 4), how does one manage curvature in discrete/combinatorial representations (Chapter 5), what novel curves arise from extrinsic directional constraints under projection (Chapter 6), and how can one measure, in an aggregate distributional sense, exactly how much distortion a given practical choice of projection introduces (Chapter 8)?

Chapter 3

Chebyshev's Variational Problem for Optimal Conformal Projections

Having established in Chapter 2 that no projection of a curved region can be an isometry, a natural engineering question arises: among all *conformal* projections of a fixed region — projections that preserve angles, hence are of the form $ds^2_{\text{plane}} = \mu(u,v)\, ds^2_{\text{sphere}}$ for a scalar scale-factor field μ — which one minimizes the worst distortion of scale?

This question was posed and answered, in a specific and elegant sense, by Pafnuty Chebyshev in a short 1856 note, decades before it received a rigorous general proof.

3.1 Setting up the variational problem

Let Ω be a fixed region on the sphere to be mapped, and let $f : \Omega \to \mathbb{R}^2$ be a conformal map (in local isothermal coordinates such a map always exists locally, by the classical theory of isothermal parameters). Conformality guarantees that the local linear scale factor

$$m(p) = \left. \frac{d(\text{planar distance})}{d(\text{spherical distance})} \right|_{\text{at } p,\ \text{any direction}} \tag{3.1}$$

is well-defined independent of direction (this is precisely what conformality buys us: the two principal stretch factors of the map's derivative coincide). Write $\sigma(p) = \log m(p)$, the logarithmic scale factor, which is 0 exactly where the map is locally isometric and departs from 0 in proportion to (logarithmic) distortion.

Chebyshev's problem: find the conformal map f minimizing

$$\max_{p \in \Omega} |\sigma(p)| \,. \tag{3.2}$$

3.2 Chebyshev's theorem

Theorem 3.1 (Chebyshev, 1856)**.** *Among all conformal projections of a simply connected region Ω of the sphere, the one minimizing the maximum logarithmic scale distortion* $\max_\Omega |\sigma|$ *is characterized by the property that σ is* constant on the boundary $\partial\Omega$.

The precise general proof (given later, and made fully rigorous by Milnor in a 1969 paper [Milnor, 1969] revisiting Chebyshev's claim [Chebyshev, 1856] with modern tools) proceeds via the following structure, which is worth presenting in full because it illustrates a genuinely elegant piece of applied harmonic analysis.

3.2.1 Why σ is harmonic

Lemma 3.2. *For a conformal map from a region of constant curvature K_0 to the (curvature-zero) plane, the logarithmic scale factor $\sigma = \log m$ satisfies the Liouville equation*

$$\Delta\sigma = -K_0\, e^{2\sigma}, \tag{3.3}$$

where Δ is the Laplacian with respect to the flat planar coordinates in which the map is expressed.

Proof idea. Writing the pulled-back spherical metric in isothermal coordinates as $ds^2 = e^{2\sigma}(du^2 + dv^2)$, the Gaussian curvature of a metric conformal to the flat one is given by the standard formula $K = -e^{-2\sigma}\Delta\sigma$. Since $K = K_0$ (constant, as we are mapping a constant-curvature spherical patch), rearranging gives the Liouville equation above. □

For *small* regions (the regime of practical cartographic interest, e.g., a single country or province, rather than a full hemisphere), σ is small and $e^{2\sigma} \approx 1$, so the Liouville equation linearizes to

$$\Delta\sigma \approx -K_0 \tag{3.4}$$

— i.e., σ satisfies a Poisson equation with constant source term $-K_0$, an inhomogeneous but linear and highly tractable PDE.

3.2.2 The extremal problem reduces to a Chebyshev-type approximation problem

The condition that $\max_\Omega |\sigma|$ be minimized, subject to σ solving a fixed Poisson equation on Ω, is structurally the same type of extremal problem that gives Chebyshev polynomials their name: among functions satisfying a fixed constraint (there, a fixed leading coefficient; here, a fixed Poisson source), minimize the maximum absolute deviation. The classical equioscillation principle underlying Chebyshev approximation has a direct geometric analogue here: the optimal σ is the solution of the Poisson equation whose boundary values are chosen so that σ is *constant on $\partial\Omega$* — pushing all of the "slack" in the maximum-principle bound for σ to the boundary, exactly where it can be absorbed uniformly, rather than allowing it to peak at an interior point.

Key Result

The Chebyshev-optimal conformal projection of a region is the one for which the scale factor is constant along the region's boundary. Interior points then automatically experience *less* distortion than the (constant) boundary value, by the maximum principle for the (elliptic) Poisson/Liouville equation governing σ: a solution of $\Delta\sigma = -K_0$ with constant boundary data attains its extrema, in the appropriate corrected sense, in a controlled way tied to the shape of $\partial\Omega$ rather than to arbitrary interior features.

3.3 A worked qualitative example: the Chebyshev–Grave projection of Russia

The projection historically associated with this theorem is the one constructed for mapping the Russian Empire, sometimes attributed jointly to Chebyshev and his student Dmitry Grave, who provided a more complete treatment in his 1896 dissertation. The construction: given the (highly non-circular, elongated) outline of the territory to be mapped, one solves the

boundary-value problem of Theorem 3.1 numerically for that specific outline, rather than using a projection — like the standard Lambert Conformal Conic — whose scale-factor behaviour is optimized for a generic band of latitude rather than tailored to the true shape of the target region.

Remark 3.3. This is a genuinely different mode of cartographic design from the classical named projections (Mercator, Lambert, stereographic, etc.), which are defined by a fixed analytic formula applicable to any region. The Chebyshev-optimal projection is instead *shape-adapted*: solving a boundary-value problem whose data (the shape of $\partial\Omega$) is the specific territory to be mapped. It anticipates, in an 1856 paper, the modern practice of designing custom "least distortion" projections for individual countries by numerical optimization — a technique now standard in national mapping agencies, executed today with finite-element PDE solvers rather than by hand.

3.4 Existence and uniqueness via the Dirichlet principle

Chebyshev's theorem *characterizes* the optimal σ by a boundary condition, but it is worth being explicit about why a solution with that boundary condition exists and is unique at all — this is not automatic for an arbitrary boundary-value problem, and the argument is a clean illustration of one of the oldest tools in the calculus of variations.

Proposition 3.4 (Existence and uniqueness of the Chebyshev-optimal σ)**.** *Let Ω be a bounded, simply connected region with sufficiently regular boundary $\partial\Omega$, and let $c \in \mathbb{R}$ be a prescribed constant. Among all functions $\sigma \in C^2(\Omega) \cap C^0(\overline{\Omega})$ satisfying $\Delta\sigma = -K_0$ in Ω and $\sigma \equiv c$ on $\partial\Omega$, there exists exactly one.*

Proof sketch. Write $\sigma = c + u$, so that u solves the homogeneous-boundary Poisson problem $\Delta u = -K_0$ in Ω, $u = 0$ on $\partial\Omega$. This is the classical Dirichlet problem for the Poisson equation. Existence follows from the Dirichlet principle: u is the unique minimizer, over all $v \in H^1_0(\Omega)$ (functions vanishing on the boundary in the appropriate Sobolev sense), of the Dirichlet energy functional

$$\mathcal{E}[v] = \int_\Omega \left(\tfrac{1}{2} |\nabla v|^2 - K_0\, v \right) dA, \tag{3.5}$$

whose Euler–Lagrange equation is exactly $\Delta v = -K_0$. The functional $\mathcal{E}$ is strictly convex and coercive on $H^1_0(\Omega)$ (a consequence of the Poincaré inequality), which guarantees both existence of a minimizer (by the direct method: take a minimizing sequence, extract a weakly convergent subsequence using coercivity, and use weak lower semicontinuity of the convex functional) and its uniqueness (strict convexity rules out two distinct minimizers). Standard elliptic regularity theory then upgrades the weak H^1 minimizer to the classical C^2 solution required by the proposition, given sufficiently smooth $\partial\Omega$. □

This existence-and-uniqueness fact is what makes Chebyshev's characterization *constructive* rather than merely a necessary condition satisfied by some unknown optimal projection: for any prescribed constant boundary value c, the Dirichlet problem has one and only one solution, and the correct value of c is then pinned down by an auxiliary global constraint (fixing the overall scale of the map, e.g. by requiring the mean value of σ over Ω to vanish, so that the map is neither globally too large nor too small on average). In practice this reduces the entire construction to a single linear elliptic boundary-value problem, which is exactly the kind of problem finite-difference and finite-element solvers were later built to handle routinely.

3.5 Milnor's rigorous treatment and a quantitative lower bound

Chebyshev's original 1856 note stated the optimality characterization without a fully rigorous justification of why the boundary-constant solution is genuinely optimal (as opposed to merely a critical point of the extremal problem); this gap was closed only in 1969, when John Milnor revisited the problem in a short paper explicitly framed as a piece of applied cartography. Milnor's contribution was twofold: a rigorous proof of Chebyshev's characterization using the maximum principle in the form sketched in Section 3.4 above, and, more strikingly, a demonstration that the unavoidable distortion is not merely nonzero but grows with the *size* of the region being mapped — larger territories cannot be conformally mapped with arbitrarily small worst-case scale error, no matter how cleverly the boundary condition is chosen, because the source term $-K_0$ in the Poisson equation acts on an area that scales with Ω itself.

Remark 3.5. This is a genuinely different and complementary statement to Theorem 2.3: Theorem 2.3 says the minimal achievable distortion is nonzero on *any* curved region, however small; Milnor's quantitative refinement says that this minimal achievable distortion, optimized over all conformal projections via Theorem 3.1, is itself an increasing function of the area of the region — so the "penalty" for mapping all of Russia conformally is unavoidably larger than the penalty for mapping a single Russian province, even under the best possible shape-adapted projection for each.

3.6 Relation to the impossibility theorem of Chapter 2

Chebyshev's theorem should be read as the constructive complement to Theorem 2.3. That theorem says: you cannot achieve $\sigma \equiv 0$ everywhere (perfect isometry) on any curved region. Chebyshev's theorem answers the natural follow-up: granted that σ cannot vanish identically, what is the *least bad* way to let it vary, within the conformal class, and how is the optimal compromise characterized? The answer — push the extremal behaviour of σ to the boundary, where the elliptic maximum principle can be used to control its interior growth — is a genuinely nontrivial piece of applied analysis that converts a philosophical impossibility into a well-posed, solvable boundary-value problem.

3.7 Summary

Chebyshev's 1856 result is a rare instance of a fully solved, sharp optimality problem in cartographic projection design predating the general theory of elliptic PDEs and conformal mapping that would later provide its rigorous justification. It demonstrates that "the least distorted map of a given territory" is not a matter of taste among a finite menu of named projections, but the solution of a specific, well-posed variational problem whose optimal solution is characterizable by a clean boundary condition (constancy of scale factor on $\partial\Omega$).

Chapter 4

The Hairy Ball Theorem and the Topology of Wind Fields

4.1 Statement of the theorem

Theorem 4.1 (Hairy Ball / Poincaré–Brouwer Theorem)**.** *There exists no continuous tangent vector field on S^2 that is everywhere nonzero. Equivalently: every continuous tangent vector field on the sphere must vanish at least once.*

The colloquial phrasing — "you cannot comb a hairy ball flat without a cowlick" — captures the content precisely: if each point of the sphere sprouts a hair lying in the tangent plane, no continuous combing pattern can avoid at least one point where the hair has no well-defined direction (a tuft, whorl, or parting).

4.2 Proof via the Poincaré–Hopf index theorem

The clean, modern proof is a corollary of a more general and equally beautiful topological result.

Theorem 4.2 (Poincaré–Hopf)**.** *Let v be a continuous vector field on a compact surface M with isolated zeros. Then the sum of the indices of the zeros of v equals the Euler characteristic of M:*

$$\sum_i \operatorname{ind}_{p_i}(v) = \chi(M). \tag{4.1}$$

For the sphere, $\chi(S^2) = 2$. Since $2 \neq 0$, the sum of indices of any vector field's zeros cannot be zero, so there must be at least one zero (a vector field with no zeros would trivially have index-sum 0). This proves Theorem 4.1 as an immediate corollary.

Sketch of Poincaré–Hopf, via triangulation. Triangulate M finely enough that v is approximately constant on each triangle, and perturb v so that its zeros occur only at vertices, are isolated, and are transverse (generic). One shows, by a careful combinatorial argument comparing the vector field's winding around the boundary of each face, edge, and vertex star, that the total index equals $V - E + F$, the alternating sum of the numbers of vertices, edges, and faces of the triangulation — which is by definition the Euler characteristic $\chi(M)$, independent of the choice of triangulation (this independence is itself a foundational fact of combinatorial topology, closely related to the discrete curvature results of Chapter 5). □

Remark 4.3. This proof is worth dwelling on because it reveals the same object — the Euler characteristic $\chi = 2$ for the sphere — appearing in three entirely different guises across this monograph: here as an index sum for vector fields, in Chapter 5 as the alternating sum $V - E + F$ governing the total angular defect of any polyhedron (Descartes's theorem, giving total defect

$4\pi = 2\pi\chi$), and again in Chapter 5 as the total spherical excess of a great-circle triangulation of the sphere itself (Girard's theorem, Section 5.3). All three are manifestations of the same topological invariant, the first continuous and combinatorial (vector-field indices on a triangulated surface), the second purely combinatorial (flat polyhedron vertices), and the third continuous and metric (curved spherical triangle areas).

4.3 Cartographic corollary: no perfectly smooth global wind map

Meteorological surface wind fields are, to an excellent approximation over short time intervals, continuous tangent vector fields on the (idealized, smooth) surface of the Earth. Applying Theorem 4.1 directly:

Corollary 4.4. *Any continuous global model of the Earth's surface wind field must contain at least one point of exactly zero wind speed (a "singularity" of the field) at any instant.*

This is not a statement of meteorological fact requiring observational data — it is a pure consequence of the sphere's topology and holds for *any* continuous vector field model, regardless of the actual physical wind pattern on a given day. Real atmospheric circulation, of course, generically contains many such points: the calm eyes of cyclones and anticyclones, and stagnation points near the poles, are the physically realized instances of the mathematically mandated zero.

Remark 4.5. It is worth being precise about what the theorem does and does not claim. It does not say that a hurricane's eye is topologically necessary in the sense of being forced to exist at a specific location, nor does it predict the number, strength, or persistence of such zeros — only that *at least one* must exist at any instant, for any continuous global wind field. A field could in principle have exactly one zero (as, for instance, the "combing" pattern that looks like flow from a single source point at one pole to a single sink at the antipodal point, familiar as the flow lines used to draw such a comb pattern on classroom globes) or many.

4.4 Interaction with map projections

A subtlety arises when one asks how this constraint appears *after* projecting the wind field onto a flat map, since flat maps are precisely where wind fields are usually visualized (weather maps, isobar/streamline charts).

Proposition 4.6. *If $\phi : S^2 \setminus \{N\} \to \mathbb{R}^2$ is any smooth diffeomorphism from the sphere minus a point (e.g., stereographic projection from the north pole N) onto the plane, and v is a continuous nonvanishing vector field on $S^2 \setminus \{N\}$, the pushed-forward field $\phi_* v$ on $\mathbb{R}^2$ is well-defined and continuous, but the theorem's force is not violated: the plane itself has $\chi(\mathbb{R}^2) = 1 \neq 0$ is not the relevant obstruction here (the plane is not compact so Poincaré–Hopf in its closed-surface form does not directly apply); instead, the constraint reappears as a statement about behaviour* at infinity *on the plane (the projected field cannot extend continuously and nonvanishingly to the point at infinity, i.e., to the one-point compactification S^2 of the plane, without violating Theorem 4.1 again) or, if the zero of the original field happened to be exactly at the removed point N, the projected planar field can appear globally nonvanishing on all of finite $\mathbb{R}^2$ while "hiding" its mandatory zero at the point at infinity, i.e. off the edge of any finite map.*

This has a genuinely practical cartographic consequence:

Key Result

A polar stereographic wind map (or any similar planar projection covering the sphere minus a point) can, in principle, display a wind field with *no visible calm point anywhere on the finite printed map*, if and only if the mandatory zero of Theorem 4.1 happens to sit exactly at the projection point removed from the plane (typically a pole). This is why global wind visualizations centred on one pole can occasionally look deceptively "zero-free": the topologically mandatory singularity has been projected to infinity, off the edge of the visible map, rather than genuinely eliminated.

4.5 Worked examples: index computation for explicit combing patterns

The Poincaré–Hopf theorem (Theorem 4.2) is often stated abstractly; it is worth making the index bookkeeping fully explicit on the two combing patterns already mentioned qualitatively in Corollary 4.4's remark, since both are standard examples in any introductory topology course and both sum correctly to $\chi(S^2) = 2$.

Example 4.7 (The single-source–single-sink field). Let v be the vector field of "meridian flow": at every point of S^2 other than the two poles N, S, let $v(p)$ be the unit tangent vector pointing along the meridian through p, directed from N towards S (this is exactly the field of arrows one would draw showing water flowing from the north pole to the south pole along lines of longitude). This field is continuous and nonvanishing on $S^2 \setminus \{N, S\}$, and has isolated zeros exactly at the two poles.

- At N, in any small disk around the pole, v points radially *outward* in the local tangent plane (away from the pole, along each meridian). A radially outward field has index $+1$ at its zero (the standard "source" index, computed by tracking that v makes exactly one full counterclockwise revolution as one circles the zero once counterclockwise).
- At S, v points radially *inward* toward the pole in the local tangent plane. A purely radial field, whether pointing inward or outward, has index $+1$ at its zero — inward-pointing ("sink") and outward-pointing ("source") radial fields are topologically indistinguishable at the level of index, since reversing every arrow of an index-$+1$ field does not change its index (the map $v \mapsto -v$ multiplies the local degree of the associated circle map by $(-1)^2 = 1$ in two dimensions).

Hence $\mathrm{ind}_N(v) + \mathrm{ind}_S(v) = 1 + 1 = 2 = \chi(S^2)$, exactly as Theorem 4.2 requires, and with exactly two zeros — the minimal number of zeros (counted as distinct points, for a generic/transverse field) that any tangent vector field on S^2 can have, since a single zero of index 2 is possible only for a non-generic (non-simple) zero, while two index-1 zeros is the generic configuration.

Example 4.8 (The "two cyclones" field). Alternatively, let w be a field with two isolated zeros at antipodal points p_1, p_2 (not necessarily the geographic poles), each a *rotational* (vortex-like) zero: in a small disk around p_i, w circulates around the zero rather than pointing radially, exactly the local picture of flow around the eye of a cyclone. A simple rotational vortex, traversed once counterclockwise around its centre, also has index $+1$ (the field itself rotates by exactly one full turn as one circles the zero once), so again $\mathrm{ind}_{p_1}(w) + \mathrm{ind}_{p_2}(w) = 2$.

Remark 4.9. These two examples make concrete exactly what Corollary 4.4 does and does not assert about real atmospheric wind fields. A physically realistic global wind field is a superposition of many local vortex-like features (individual cyclones and anticyclones, each locally

resembling Example 4.8) rather than a single clean source–sink pattern (Example 4.7), and the theorem is entirely agnostic about *how* the mandatory total index of 2 is distributed among however many zeros the real field happens to have — it could be two simple index-1 vortices as in a quiet week, or dozens of simple vortices of mixed sign whose indices still sum to exactly 2 (a field can certainly have zeros of negative index, such as saddle-type zeros with index -1, provided enough compensating positive-index zeros exist elsewhere to keep the total at 2; a hurricane season with three positive vortices and one saddle-like zero of index -1 would satisfy $3(+1) + (-1) = 2$ just as validly as two simple vortices would).

4.6 A note on the two-dimensional torus, for contrast

It is instructive to note what makes the sphere special here. The torus T^2 has $\chi(T^2) = 0$, and indeed nonvanishing continuous vector fields on the torus exist trivially (e.g., the constant field in one of the two natural periodic directions). This is precisely why, for a hypothetical planet shaped like a torus, no topological obstruction of this kind would force the existence of a calm point in a smooth model wind field — the Hairy Ball obstruction is a genuinely spherical (more precisely, genus-zero) phenomenon, tied to $\chi = 2$, not a universal feature of "any closed surface with a wind field."

4.7 Summary

The Hairy Ball Theorem converts a piece of pure algebraic topology — the non-vanishing of the Euler characteristic of the sphere — into an unconditional constraint on every continuous global vector field defined on it, wind fields included. Unlike the results of Chapters 2 and 3, which concern the metric (distance/angle) structure of the sphere, this result concerns only its topology, and would hold even if the Earth's surface metric were altered while its genus-zero, sphere-like topology were preserved.

Chapter 5

Discrete Curvature: Angular Defect and Polyhedral Projections

Chapters 2–4 treat the sphere as a smooth object. Many practical cartographic constructions — most famously Buckminster Fuller's Dymaxion map, but also the family of "orange-peel" (interrupted, gored) projections and modern discrete-global-grid systems used in geographic information science — instead approximate the sphere by a polyhedron and unfold its faces onto the plane. This chapter develops the discrete analogue of Gaussian curvature that governs such constructions.

5.1 Angular defect

Definition 5.1. Let P be a convex polyhedron and v one of its vertices, at which faces meet contributing interior face-angles $\alpha_1, \dots, \alpha_k$ (one per face incident to v). The *angular defect* at v is

$$\delta(v) = 2\pi - \sum_{i=1}^{k} \alpha_i. \tag{5.1}$$

Angular defect is the combinatorial, vertex-localized stand-in for Gaussian curvature: on a flat plane, six equilateral triangles (or four squares, or three hexagons) meet at a point with angles summing to exactly 2π, giving zero defect; on a polyhedron, the angles around a vertex sum to *less* than 2π (for a convex vertex), giving positive defect, exactly mirroring the fact that positive Gaussian curvature is associated with angle sums around a point that fall short of the flat value.

5.2 Descartes's theorem

Theorem 5.2 (Descartes, *ca.* 1630; rediscovered/published by Legendre, 1794)**.** *For any convex polyhedron P topologically equivalent to a sphere (genus 0), the sum of the angular defects over all vertices equals 4π:*

$$\sum_{v} \delta(v) = 4\pi. \tag{5.2}$$

Proof. Let V, E, F be the numbers of vertices, edges, and faces. Each face with n_i sides has interior angle sum $(n_i - 2)\pi$ (for a planar polygon), so the total angle sum over all faces is

$$\sum_{\text{faces}} (n_i - 2)\pi = \pi \left(\sum_{i} n_i \right) - 2\pi F = \pi(2E) - 2\pi F = 2\pi E - 2\pi F, \tag{5.3}$$

using $\sum_i n_i = 2E$ (each edge borders exactly two faces). This total angle sum is exactly the sum, over all vertices, of the angle sums $\sum_i \alpha_i$ contributed at that vertex (every face-angle at

every vertex is counted once). Hence

$$\sum_v \sum_i \alpha_i = 2\pi E - 2\pi F. \tag{5.4}$$

Now sum the defect definition over all vertices:

$$\sum_v \delta(v) = \sum_v \left(2\pi - \sum_i \alpha_i\right) = 2\pi V - (2\pi E - 2\pi F) = 2\pi(V - E + F) = 2\pi\chi(P). \tag{5.5}$$

For any polyhedron homeomorphic to the sphere, Euler's formula gives $\chi(P) = V - E + F = 2$, so $\sum_v \delta(v) = 4\pi$. □

Remark 5.3. This is a purely combinatorial argument — no calculus, no limits, no smoothness — and yet it produces exactly the same constant, 4π, that appears as $\int_{S^2} K\, dA$ for the smooth sphere under the (continuous) Gauss–Bonnet theorem: $\int_{S^2} K\, dA = 4\pi$ for a sphere of any radius R, since $K = 1/R^2$ and the surface area is $4\pi R^2$. Descartes's theorem is thus a genuinely discrete, combinatorial Gauss–Bonnet theorem, predating the general smooth Gauss–Bonnet theorem (usually credited to Bonnet's 1848 formulation [Bonnet, 1848] building on Gauss) by roughly two centuries, and holding exactly (not merely as an approximation) for every genus-zero polyhedron regardless of its specific shape.

Key Result

The total angular defect of *any* convex polyhedron topologically equivalent to a sphere — whether a cube, a dodecahedron, a geodesic dome with thousands of faces, or an irregular lump with straight edges — is always exactly 4π radians (equivalently, 720°). This total is a topological invariant, completely independent of the polyhedron's particular geometry; only its *distribution* across the vertices depends on shape.

5.3 Girard's theorem: spherical excess as the continuous counterpart

Descartes's theorem (Theorem 5.2) concentrates curvature at the *vertices* of a flat, straight-edged polyhedron. It is natural to ask what the same accounting looks like if, instead of flattening the faces, one keeps every face a genuine curved patch of the sphere itself, bounded by great-circle arcs rather than straight edges — i.e., a *spherical* polygon rather than a planar one. The classical answer, predating Descartes's own result by roughly a decade, is due to the French mathematician Albert Girard [Girard, 1629].

Theorem 5.4 (Girard, 1629)**.** *Let T be a spherical triangle on a sphere of radius R, bounded by three great-circle arcs, with interior angles A, B, C (measured in radians, at the three vertices where the arcs meet). Then the area of T is*

$$\text{area}(T) = R^2\,(A + B + C - \pi). \tag{5.6}$$

The quantity $E(T) = A + B + C - \pi$ is called the spherical excess *of T.*

Proof. A *lune* of angle α is the region of the sphere bounded by two half great-circles sharing both endpoints (a pair of antipodal points) and meeting there at angle α. Since a lune of angle 2π (i.e., the full sphere cut along a single great circle and reassembled) has area $4\pi R^2$, and lune area scales linearly with α, a lune of angle α has area $2\alpha R^2$.

Extend the three great circles carrying the sides of T to full great circles. Three great circles in general position partition the sphere into eight spherical triangles, occurring in four antipodal

pairs; antipodal reflection is an isometry of the sphere, so each pair consists of two triangles of equal area. Write T for the given triangle (with angles A, B, C and area written simply as T by abuse of notation for its area), and let T_A, T_B, T_C denote the three triangles among the remaining seven that each share exactly one vertex of T (namely A, B, C respectively) and lie on the far side of the opposite side's great circle, so that the lune of angle A at vertex A is exactly the union $T \cup T_A$ (and similarly for B, C). Since the eight triangles form four antipodal pairs each of equal area, and T, T_A, T_B, T_C together with their four antipodal partners exhaust all eight,

$$2\,(T + T_A + T_B + T_C) = 4\pi R^2 \quad \Longrightarrow \quad T_A + T_B + T_C = 2\pi R^2 - T. \tag{5.7}$$

On the other hand, summing the three lune areas directly,

$$\begin{aligned} 2AR^2 + 2BR^2 + 2CR^2 &= (T + T_A) + (T + T_B) + (T + T_C) \\ &= 3T + (T_A + T_B + T_C) = 3T + 2\pi R^2 - T = 2T + 2\pi R^2. \end{aligned} \tag{5.8}$$

Solving for T:

$$T = R^2(A + B + C) - \pi R^2 = R^2(A + B + C - \pi), \tag{5.9}$$

which is (5.6). □

Corollary 5.5 (Spherical excess of a spherical n-gon)**.** *Let Q be a spherical polygon with n vertices, bounded by n great-circle arcs, with interior angles $\alpha_1, \ldots, \alpha_n$, and suppose Q can be triangulated into $n - 2$ spherical triangles by great-circle diagonals (as any convex, and more generally any simple, spherical polygon can). Then*

$$\text{area}(Q) = R^2 \left(\sum_{i=1}^{n} \alpha_i - (n-2)\pi \right). \tag{5.10}$$

Proof. Summing (5.6) over the $n - 2$ triangles of the triangulation: each triangle contributes its own three angles, and at every diagonal the two angles on either side of it sum to a full angle at that point which is entirely internal to Q's interior, so the grand total of all triangle-angle contributions equals $\sum_i \alpha_i$ (the polygon's own vertex angles) plus, at each of the $n - 3$ interior diagonal endpoints where more than two triangle corners might meet, angles that are exhaustively accounted for by the same sum; a direct induction on the number of triangles (adding one triangle at a time by cutting off one ear of the polygon) shows the total area is $R^2 \sum_i \alpha_i - (n-2)\pi R^2$, since each added triangle contributes $-\pi R^2$ to the constant term and each ear-cutting step preserves the sum of boundary angles exactly. The base case $n = 3$ is Theorem 5.4 itself. □

Remark 5.6 (Girard's theorem as the exact continuous analogue of Descartes's theorem). Theorem 5.2 and Theorem 5.4 are, in a precise sense, the same statement applied to two different combinatorial objects sharing the same underlying sphere. Descartes's theorem concentrates the sphere's total curvature budget 4π at the finitely many *vertices* of a flat polyhedron (angular defect, a purely combinatorial quantity with no reference to any metric on the faces themselves, since the faces are literally flat). Girard's theorem instead distributes that same budget continuously, in a genuinely metric (area-dependent) way, over the *faces* of a spherical (not flattened) triangulation: each spherical triangle in the tiling carries a nonzero excess $E(T) = A + B + C - \pi$ exactly in proportion to its own area, per (5.6), and summing $E(T)$ over every triangle of *any* triangulation of the whole sphere into spherical triangles returns the same total,

$$\sum_{\text{triangles } T} E(T) = \frac{1}{R^2} \sum_{T} \text{area}(T) = \frac{4\pi R^2}{R^2} = 4\pi, \tag{5.11}$$

regardless of how many triangles the triangulation uses or how they are shaped — a face-based invariance exactly parallel to the vertex-based invariance of Theorem 5.2. Both are, in turn, instances of the single Gauss–Bonnet identity $\int_M K\,dA = 2\pi\chi(M)$ already invoked in the remark following Theorem 5.2: Descartes's $4\pi = 2\pi\chi(P)$ is its purely combinatorial (vertex-supported) specialization to a flat polyhedron, while Girard's theorem, summed over a triangulation, is its exact specialization to the smooth, constantly-curved sphere itself, with the curvature $K = 1/R^2$ now spread continuously across each triangle's interior rather than concentrated at isolated points. This gives a third manifestation, alongside the vector-field index sum of Chapter 4 and the angular defect sum of Theorem 5.2, of the single topological invariant $\chi(S^2) = 2$.

Example 5.7 (Spherical excess of an icosahedral face). Return to the icosahedron of Section 5.4 (used there for the Dymaxion map), but now consider its 20 vertices connected by great-circle arcs on the circumscribed sphere, rather than by the straight edges of the flat polyhedron — i.e., the spherical triangulation of S^2 combinatorially dual to the flat icosahedron's vertex/edge/-face structure ($V = 12$, $E = 30$, $F = 20$ carried over unchanged, since the triangulation has the same combinatorics, only with curved rather than straight edges). By the regular icosahedron's five-fold symmetry, exactly 5 congruent spherical triangles meet at each of the 12 vertices, and since a vertex of a smooth sphere has a full 2π of angle around it (unlike a polyhedron vertex, where the angle sum falls short of 2π by the angular defect of Definition 5.1), the five equal spherical-triangle angles meeting there must satisfy

$$5\,\alpha = 2\pi \quad \Longrightarrow \quad \alpha = \frac{2\pi}{5} = 72^\circ. \tag{5.12}$$

Each of the 20 congruent spherical triangles is equilateral with all three angles equal to $\alpha = 2\pi/5$, so by Theorem 5.4 each has area

$$\text{area}(T) = R^2\left(3\cdot\frac{2\pi}{5} - \pi\right) = R^2\left(\frac{6\pi}{5} - \pi\right) = \frac{\pi R^2}{5}. \tag{5.13}$$

Summing over all 20 faces gives $20\cdot\pi R^2/5 = 4\pi R^2$, exactly the total surface area of the sphere, confirming the computation is consistent to machine precision with no free parameters: the vertex angle $\alpha = 2\pi/5$ was fixed entirely by the icosahedron's five-fold vertex symmetry, and Girard's formula then reproduces the correct total area with no further adjustment. This is the exact spherical-triangulation counterpart to Section 5.5's flat-polyhedron computation, where the same 12 vertices instead carry a flat angular defect of $4\pi/12 = \pi/3$ each (Section 5.4): the two pictures — flat angular defect at vertices, spherical excess across faces — are complementary discrete and continuous accountings of the same sphere, differing because a flat polyhedron face and its corresponding spherical triangle are related by the gnomonic projection (Theorem 7.2), which is neither conformal nor angle-preserving, so the flat face's 60° equilateral angles and the spherical triangle's 72° equilateral angles are genuinely different numbers, each correct for its own (flat vs. curved) geometric object.

5.4 Polyhedral map projections

A polyhedral (or "gored"/interrupted) map projection replaces the smooth sphere with an inscribed or circumscribed polyhedron, projects the sphere's surface (typically via gnomonic or other elementary projections) onto each planar face separately, and then unfolds the polyhedron's faces into a connected planar net.

5.4.1 The Dymaxion map

Buckminster Fuller’s 1943 Dymaxion (Airocean World) map [Fuller, 1943] uses an icosahedron ($V = 12$, $E = 30$, $F = 20$, each face an equilateral triangle) as the reference polyhedron. Each of the sphere’s twenty triangular regions (bounded by great-circle arcs corresponding to the icosahedron’s edges projected outward onto the sphere) is mapped, with comparatively low internal distortion because each triangular patch subtends a relatively small angular extent of the sphere, onto the corresponding flat icosahedral face, and the twenty faces are then unfolded into a net that can be arranged in various ways (Fuller favoured an arrangement emphasizing the continuity of the world’s oceans or landmasses, depending on the specific unfolding chosen).

Remark 5.8. By Theorem 5.2, the twelve vertices of the icosahedron used as the reference solid carry, between them, a total angular defect of exactly 4π; for the regular icosahedron this is distributed *equally*, $\delta(v) = 4\pi/12 = \pi/3$ at each of the twelve (identical, by symmetry) vertices. This uniform distribution of defect is precisely why the regular icosahedron is a comparatively good choice of reference solid for low, evenly spread distortion: no single vertex region is forced to absorb an outsized share of the sphere’s total curvature budget, unlike, say, a cube (only $V = 8$ vertices, so $\delta(v) = 4\pi/8 = \pi/2$ per vertex if the defect is assumed evenly spread by symmetry) or an irregular polyhedron whose defect could concentrate arbitrarily at a small number of vertices.

5.4.2 Why more faces (finer polyhedra) reduce distortion, quantitatively

As the reference polyhedron is refined — moving from an icosahedron (20 faces) toward a geodesic subdivision with hundreds or thousands of small triangular faces — the angular extent of the sphere subtended by each individual face shrinks. Since each face is still projected by an ordinary (non-isometric, per Theorem 2.3) planar projection, some distortion always remains within each face, but the *magnitude* of that within-face distortion is controlled by the extent of the region being projected: a standard estimate for, e.g., a gnomonic projection of a spherical cap of angular radius θ shows the local scale factor deviates from 1 by an amount of order θ^2 for small θ (a second-order effect, consistent with the fact that at $\theta = 0$, i.e. at the tangent point, the projection is exactly isometric to first order). Halving the angular size of each face therefore reduces the leading-order within-face distortion by a factor of approximately four, at the cost of increasing the number of face-boundary seams (along which the map is necessarily discontinuous or must be specially blended) by a comparable factor. Polyhedral projection design is thus a genuine trade-off, quantitatively controlled by Theorem 5.2, between few large low-fidelity faces and many small high-fidelity faces with more numerous seams.

5.5 Geodesic subdivision: making “finer polyhedra” precise

The qualitative trade-off described in the previous section — more, smaller faces reduce within-face distortion at the cost of more seams — can be made exact using the standard combinatorics of *geodesic subdivision*, the construction underlying geodesic domes and the finer discrete-global-grid systems used in practice.

Definition 5.9 (Class-I geodesic subdivision of frequency n)**.** Starting from the icosahedron ($V_0 = 12$, $E_0 = 30$, $F_0 = 20$), subdivide each of the 20 triangular faces into a regular grid of n^2 smaller triangles (frequency n: each original edge is divided into n equal segments), and radially project every resulting vertex outward onto the circumscribed sphere.

Proposition 5.10. *The frequency-n geodesic subdivision of the icosahedron has*

$$F = 20n^2, \qquad E = 30n^2, \qquad V = 10n^2 + 2. \tag{5.14}$$

Proof. Each of the 20 original faces is divided into n^2 small triangles, giving $F = 20n^2$ directly. Each small triangle has 3 edges and each edge is shared by exactly 2 triangles, so $E = 3F/2 = 30n^2$. Euler's formula $V - E + F = 2$ (valid since the subdivided, sphere-projected solid remains topologically a sphere) then gives $V = 2 - F + E = 2 - 20n^2 + 30n^2 = 10n^2 + 2$. □

Remark 5.11. Setting $n = 1$ recovers the original icosahedron exactly ($F = 20, E = 30, V = 12$), as it must. Theorem 5.2 continues to hold at every frequency: $\sum_v \delta(v) = 4\pi$ regardless of n, since the subdivided solid is still a genus-zero convex polyhedron. What changes with n is only the *distribution* of that fixed 4π budget: it is spread over $V = 10n^2 + 2$ vertices instead of 12, so the *typical* (though not exactly uniform, since the original 12 icosahedral vertices remain geometrically distinguished) angular defect per vertex scales like $4\pi/(10n^2+2) = O(n^{-2})$ as n grows. This is the combinatorial counterpart of the $O(\theta^2)$ scale-factor estimate already invoked in the previous section for the continuous gnomonic projection of a spherical cap of angular radius $\theta \sim 1/n$: both express, in their respective discrete and continuous languages, the same fact that curvature concentrated over a shrinking angular extent produces a proportionally shrinking local geometric effect.

Example 5.12. A frequency-4 subdivision of the icosahedron ($n = 4$) has $F = 20(16) = 320$ triangular faces, $E = 30(16) = 480$ edges, and $V = 10(16) + 2 = 162$ vertices — a scale of refinement typical of geodesic dome architecture and of coarse discrete global grids used for climate-model output, still small enough that $\sum_v \delta(v) = 4\pi$ is carried by only 162 vertices rather than the 612-point global sampling lattice used computationally in Chapter 8.

5.6 A discrete Gauss–Bonnet with boundary

For non-closed polyhedral surfaces (e.g., a single unfolded net with a boundary, as used in practice once the polyhedron is cut open along some edges to be laid flat), the discrete Gauss–Bonnet theorem takes a form with a boundary term, directly paralleling the smooth Gauss–Bonnet theorem with geodesic boundary curvature:

$$\sum_{v \in \text{interior}} \delta(v) + \sum_{v \in \partial} \delta_\partial(v) = 2\pi\chi(P), \tag{5.15}$$

where $\delta_\partial(v) = \pi - \sum_i \alpha_i$ is the appropriately modified boundary defect (a flat boundary vertex has angle sum π, not 2π), and $\chi(P) = 1$ for a disk-like (simply connected, once-cut-open) net. This is the discrete combinatorial ancestor of the general smooth Gauss–Bonnet formula $\int_M K\,dA + \int_{\partial M} \kappa_g\,ds = 2\pi\chi(M)$, and it is precisely the identity that guarantees a polyhedral net, once fully unfolded and laid flat, has total angle "slack" at its vertices and boundary summing in a way that is fixed by topology alone, regardless of how the cutting and unfolding is performed.

5.7 Summary

Angular defect is the combinatorial shadow of Gaussian curvature, and Descartes's theorem — a purely combinatorial identity provable with nothing beyond Euler's formula and elementary polygon angle-sum facts — is the discrete Gauss–Bonnet theorem, fixing the total curvature

"budget" of any genus-zero polyhedron at exactly 4π regardless of shape. Girard's theorem (Theorem 5.4) supplies the exact continuous counterpart, distributing that same 4π budget as spherical excess across the faces of a curved (great-circle-bounded) triangulation rather than as angular defect at the vertices of a flattened one, and Example 5.7 verifies the two accountings side by side on the same icosahedral combinatorics. Together, these results are the mathematical bedrock underlying every polyhedral map projection, from Fuller's Dymaxion map to modern discrete global grid systems used in climate modelling and geographic databases, and they explain quantitatively why refining the reference polyhedron (more, smaller faces) trades reduced within-face distortion against a proliferation of unavoidable seams.

Chapter 6

Loxodromes as Logarithmic Spirals of Finite Length

6.1 Definition

Definition 6.1. A *loxodrome* (or *rhumb line*) on the sphere is a curve that crosses every meridian at a constant angle β (the compass bearing), where $\beta \neq 0, \pi/2$.

Loxodromes are of direct historical and practical navigational importance: a ship or aircraft holding a constant compass bearing traces a loxodrome, not (in general) the shorter great-circle route between its endpoints. Understanding their geometry precisely is a matter of substantial historical consequence in the age of sail.

6.2 Parametrization and the differential equation

Let (φ, λ) denote latitude and longitude. Along a curve making constant angle β with meridians, an infinitesimal displacement $(d\varphi, d\lambda)$ on the sphere of radius R has north-south arc length $R\,d\varphi$ and east-west arc length $R\cos\varphi\,d\lambda$, and the bearing condition $\tan\beta = \dfrac{\text{east-west component}}{\text{north-south component}}$ gives

$$\tan\beta = \frac{R\cos\varphi\,d\lambda}{R\,d\varphi} = \cos\varphi\,\frac{d\lambda}{d\varphi}. \tag{6.1}$$

Separating variables:

$$d\lambda = \tan\beta\sec\varphi\,d\varphi \quad\Longrightarrow\quad \lambda(\varphi) = \lambda_0 + \tan\beta\int_{\varphi_0}^{\varphi}\sec\varphi'\,d\varphi'. \tag{6.2}$$

The integral of $\sec\varphi$ is the classical Gudermannian-related integral:

$$\int\sec\varphi\,d\varphi = \ln\left|\tan\left(\frac{\varphi}{2}+\frac{\pi}{4}\right)\right| + C, \tag{6.3}$$

giving the explicit loxodrome equation

$$\lambda(\varphi) = \lambda_0 + \tan\beta\,\left[\ln\left|\tan\left(\frac{\varphi}{2}+\frac{\pi}{4}\right)\right| - \ln\left|\tan\left(\frac{\varphi_0}{2}+\frac{\pi}{4}\right)\right|\right]. \tag{6.4}$$

Remark 6.2. The quantity $\ln\tan(\varphi/2+\pi/4)$ is exactly the (unscaled) Mercator y-coordinate as a function of latitude (compare Chapter 8, Section 8.2.3). This is not a coincidence: it is precisely the reason the Mercator projection was invented. Under Mercator's projection, λ is mapped linearly to the plane's x-coordinate and $\ln\tan(\varphi/2+\pi/4)$ linearly to the y-coordinate, so the loxodrome equation above becomes exactly $y = \tan\beta\cdot x + \text{const}$: a *straight line* of constant slope $\tan\beta$. This is the entire reason for Mercator's 1569 construction: to make loxodromes — the curves a navigator actually sails — appear as straight lines that can be drawn with a ruler on the chart.

6.3 The loxodrome as a logarithmic spiral

While the Mercator picture (straight line) is the practically important one, the loxodrome's shape *on the sphere itself*, viewed for instance in stereographic projection centred on a pole, is a logarithmic spiral.

Proposition 6.3. *Under stereographic projection from a pole, a loxodrome making constant bearing angle β with meridians maps to a logarithmic spiral $r(\theta) = r_0\, e^{-\theta \cot\beta}$ (in polar coordinates (r, θ) centred at the image of the pole), winding infinitely many times around the origin as it approaches the pole.*

Proof sketch. Stereographic projection from a pole is conformal and maps meridians to straight rays from the origin, and parallels to concentric circles. A curve crossing every ray (meridian-image) at the same constant angle β (angle preservation under conformal maps) is, by the classical characterization of curves of constant angle to a pencil of rays through a point, exactly a logarithmic (equiangular) spiral $r = r_0 e^{-\theta \cot\beta}$. □

This equiangular/logarithmic spiral property is the same mathematical curve that appears famously in the growth pattern of the nautilus shell and in certain phyllotactic (plant growth) spirals — in all these cases arising from the same underlying principle: a curve maintaining constant angle to a family of rays or a self-similar growth process.

6.4 Infinite winding, finite length

Here is the genuinely striking fact, easy to underappreciate without doing the integral explicitly.

Theorem 6.4. *A loxodrome making a fixed non-orthogonal, non-meridional bearing β winds around each pole* infinitely many times *as it approaches it (the longitude $\lambda(\varphi) \to \pm\infty$ as $\varphi \to \pm\pi/2$), yet the total arc length of the loxodrome from any starting latitude φ_0 to the pole is* finite.

Proof. The infinite winding is immediate from the formula for $\lambda(\varphi)$: as $\varphi \to \pi/2^-$, $\tan(\varphi/2 + \pi/4) \to \infty$, so $\ln\tan(\varphi/2 + \pi/4) \to \infty$, and hence $\lambda(\varphi) \to \pm\infty$ (sign depending on the sign of $\tan\beta$) — the curve's longitude coordinate increases without bound, meaning it crosses every meridian infinitely often, i.e., spirals around the pole infinitely many times.

For the arc length: an infinitesimal element of arc length along the loxodrome, using $\tan\beta = \cos\varphi\, d\lambda/d\varphi$ to eliminate $d\lambda$, is

$$ds = \sqrt{(R\,d\varphi)^2 + (R\cos\varphi\, d\lambda)^2} = R\,d\varphi\sqrt{1 + \tan^2\beta} = R\sec\beta\, d\varphi \tag{6.5}$$

(using $\cos\varphi\, d\lambda = \tan\beta\, d\varphi$ from the defining relation). Crucially, ds is proportional to $d\varphi$ alone, with a *constant* proportionality factor $R\sec\beta$ — the φ-dependence has entirely cancelled out. The total arc length from φ_0 to the pole $\varphi = \pi/2$ is therefore simply

$$L = \int_{\varphi_0}^{\pi/2} R\sec\beta\, d\varphi = R\sec\beta\left(\frac{\pi}{2} - \varphi_0\right), \tag{6.6}$$

manifestly finite for any $\varphi_0 > -\pi/2$ and any fixed $\beta \neq \pi/2$. □

Key Result

A loxodrome spirals around each pole infinitely many times, yet covers only a finite total distance in doing so — a vivid, elementary example of a curve with infinite angular winding and finite arc length, provable with nothing beyond a single change-of-variables integral. The apparent paradox (infinitely many loops in finite length) is resolved by the fact that the loops shrink geometrically: because $\lambda(\varphi)$ grows only *logarithmically* as $\varphi \to \pi/2$, while $ds/d\varphi$ stays exactly constant, successive loops around the pole occur over exponentially shrinking intervals of φ near the pole, and the geometric shrinkage is fast enough that the total length — an ordinary convergent improper integral, not a subtler summation — stays finite.

6.5 A remark on series convergence

The finiteness of L can also be seen, in a form closer to a convergent series, by summing the length contributed by each successive full loop around the pole. Since each loop corresponds to λ advancing by 2π, and $\lambda(\varphi) - \lambda_0 = \tan\beta \cdot [\ln\tan(\varphi/2 + \pi/4) - \ln\tan(\varphi_0/2 + \pi/4)]$ grows logarithmically, the latitude values φ_n at which successive loops complete satisfy $\tan(\varphi_n/2 + \pi/4) \sim C \cdot e^{2\pi n \cot\beta}$ for a constant C — i.e., φ_n approaches $\pi/2$ at a rate governed by a geometric (exponential) sequence in n. Since (from the proof above) the arc length *per unit* φ is constant, and the φ-intervals between successive loops shrink geometrically, the length of the n-th loop shrinks geometrically in n, and the total length is the sum of a convergent geometric-type series — structurally the same convergence mechanism as $\sum_n r^n$ for $|r| < 1$, here arising naturally from the sphere's geometry rather than being imposed as a modelling choice.

6.6 The general rhumb-line distance formula, and a worked numerical example

The finite-length computation of Theorem 6.4 was carried out for the special case of a loxodrome running all the way to a pole, but the same derivation applies verbatim, and more usefully for ordinary navigation, between any two latitudes $\varphi_0 < \varphi_1$ (neither of which need be a pole).

Corollary 6.5 (Rhumb-line distance between two latitudes)**.** *The arc length of the loxodrome of bearing β between latitudes φ_0 and φ_1 is*

$$L(\varphi_0, \varphi_1, \beta) = R \sec\beta\,(\varphi_1 - \varphi_0), \tag{6.7}$$

independent of the change in longitude, since (as shown in the proof of Theorem 6.4) $ds = R\sec\beta\,d\varphi$ regardless of how many times the curve has wound around a pole in reaching φ_1.

Example 6.6. Take a rhumb-line course of bearing $\beta = 50^\circ$ from the ship's compass, starting at latitude $\varphi_0 = 10^\circ$N and ending at latitude $\varphi_1 = 50^\circ$N, on the IUGG mean-radius spherical Earth model $R = 6371.0088\,\text{km}$ used throughout this monograph (Chapter 8). By Corollary 6.5,

$$L = R\sec(50^\circ)\left(\frac{50\pi}{180} - \frac{10\pi}{180}\right) \approx 6371.0088 \times 1.5557 \times 0.6981\,\text{km} \approx 6919.55\,\text{km}, \tag{6.8}$$

and the corresponding change in longitude, from Eq. (6.4) (the loxodrome equation derived above),

$$\Delta\lambda = \tan(50^\circ)\left[\ln\tan\left(\frac{50^\circ}{2}+45^\circ\right) - \ln\tan\left(\frac{10^\circ}{2}+45^\circ\right)\right] \approx 0.99542 \text{ rad} \approx 57.03^\circ. \tag{6.9}$$

This single example already illustrates the practical value of Corollary 6.5: the distance sailed depends only on the latitude change and the bearing, via the clean factor $R\sec\beta$, while all of the geometric complexity of the sphere's convergent meridians is absorbed into the (still exact, but transcendental) longitude formula — exactly the asymmetry that motivated separating the two computations historically, with tables of $\ln\tan(\varphi/2+\pi/4)$ (the "meridional parts" of traditional navigation manuals) precomputed for every minute of latitude so that navigators could obtain $\Delta\lambda$ by simple subtraction and multiplication, without evaluating a logarithm by hand at sea.

Remark 6.7. Before Mercator's meridional-parts tables became standard, navigators often used the cruder *middle-latitude sailing* approximation, which replaces the exact $\sec\varphi$ weighting inside the integral of Eq. (6.4) with a single constant evaluated at the mean latitude $\varphi_m = (\varphi_0+\varphi_1)/2$, i.e. approximating $\Delta\lambda \approx \tan\beta\,(\varphi_1-\varphi_0)\sec\varphi_m$ instead of the exact logarithmic expression. This approximation is adequate over short passages at moderate latitude, where $\sec\varphi$ varies slowly, but degrades increasingly as either the latitude span or the absolute latitude grows, since $\sec\varphi$ is a convex, rapidly increasing function near the poles — precisely the regime where the exact formula of Corollary 6.5 and the logarithmic longitude formula are needed instead.

6.7 Rhumb line versus great circle: a quantitative comparison

The rhumb line of Example 6.6 is convenient to sail (constant compass bearing) but is not, in general, the shortest path between its endpoints: the shortest path on the sphere is always the great-circle (orthodromic) arc, whose length is given by the haversine formula already used throughout Chapter 8 (Eq. (8.2) there). It is instructive to compute both distances for the same pair of endpoints and compare them directly.

Example 6.8. Take the same rhumb line as Example 6.6: constant bearing $\beta = 50^\circ$ from $(\varphi_0,\lambda_0) = (10^\circ\text{N}, 0^\circ)$ to $\varphi_1 = 50^\circ\text{N}$, which by Eq. (6.4) arrives at longitude $\lambda_1 = \lambda_0 + \Delta\lambda \approx 57.03^\circ\text{E}$. The rhumb-line distance, from Corollary 6.5, is $L_{\text{rhumb}} \approx 6919.55$ km. The great-circle distance between the same two endpoints $(10^\circ\text{N}, 0^\circ)$ and $(50^\circ\text{N}, 57.03^\circ\text{E})$, computed with the haversine formula on the same spherical Earth model $R = 6371.0088$ km, is

$$L_{\text{gc}} \approx 6836.16 \text{ km}. \tag{6.10}$$

The rhumb line is therefore longer by $L_{\text{rhumb}} - L_{\text{gc}} \approx 83.39$ km, an excess of about 1.22% over the great-circle distance ($L_{\text{rhumb}}/L_{\text{gc}} \approx 1.0122$). The initial bearing of the great-circle route at the starting point, by the standard initial-bearing formula, is approximately 37.86° — markedly different from the rhumb line's constant 50° bearing — and the great-circle bearing would have to change continuously along the route (curving initially further from the pole-ward direction and steepening as the path approaches its northernmost extent) to stay on the shorter path, which is exactly the navigational inconvenience the rhumb line is designed to avoid at the cost of the modest distance penalty just computed.

Remark 6.9. This 1.22% penalty is representative rather than universal: the fractional excess of rhumb-line over great-circle distance grows with the angular separation of the endpoints and, for a fixed separation, with how far the bearing β departs from the great-circle's own

initial bearing (a rhumb line running due north–south, $\beta = 0$, coincides exactly with a meridian, which is itself a great circle, giving zero excess; the excess grows as β departs from the locally great-circle-matching bearing). Example 6.8 illustrates, with one concrete and fully reproducible pair of numbers, the general principle already implicit in Theorem 6.4: the rhumb line's convenience (constant bearing, hence a straight line on a Mercator chart) is purchased at a real but usually modest cost in distance sailed, and the size of that cost is exactly quantifiable by comparing Corollary 6.5 against the haversine great-circle formula for the same endpoints.

6.8 Summary

The loxodrome is the paradigm example, in this monograph, of a curve whose striking mathematical behaviour (a logarithmic spiral of infinite angular winding but finite length) is a direct consequence of an *extrinsic*, direction-dependent definition (constant bearing relative to a fixed reference direction, the meridians converging at the poles) rather than of any intrinsic curvature property of the sphere. It is also the historical motivation for the single most famous of all map projections, Mercator's, whose entire defining property (loxodromes map to straight lines) is a direct transcription of the differential equation derived in Chapter 6 above.

Chapter 7

The Gnomonic and Stereographic Projections: Two Classical Circle-Preserving Constructions

The preceding chapters have already invoked two named projections without pausing to establish their defining geometric properties in full: the gnomonic projection, used implicitly in Chapter 5 to map each triangular face of a reference polyhedron such as the icosahedron, and the stereographic projection, used in Chapter 6 to display the loxodrome as a logarithmic spiral. Both belong to the family of *azimuthal* (or *zenithal*) projections — projections defined by a single tangent point (or, equivalently, a plane tangent to the sphere at one point), under which every point's image lies at a distance from the centre, and in a direction from the centre, that depend only on the point's angular distance c and bearing from the tangent point. This chapter isolates the two members of that family with the sharpest classical theorems attached to them, and proves both theorems in full.

7.1 The azimuthal projection family

Let (φ_1, λ_0) be the tangent point, and let c denote the great-circle angular distance from the tangent point to a general point (φ, λ), given by

$$\cos c = \sin\varphi_1 \sin\varphi + \cos\varphi_1 \cos\varphi \cos(\lambda - \lambda_0). \tag{7.1}$$

Every azimuthal projection can be written in the common form (Snyder, *Map Projections: A Working Manual*, 1987 [Snyder, 1987], Ch. 5)

$$x = R\,k' \cos\varphi \sin(\lambda - \lambda_0), \qquad y = R\,k' \left[\cos\varphi_1 \sin\varphi - \sin\varphi_1 \cos\varphi \cos(\lambda - \lambda_0)\right], \tag{7.2}$$

where the single scalar function $k' = k'(c)$ — constant on each circle of fixed angular distance c from the tangent point — distinguishes the individual named projections within the family:

7.2 The gnomonic projection: great circles as straight lines

Definition 7.1 (Gnomonic projection)**.** The gnomonic projection maps a point p of the sphere to the point where the straight ray from the sphere's centre O through p meets the plane tangent to the sphere at (φ_1, λ_0). It is defined only for $c < \pi/2$ (the open hemisphere centred at the tangent point), since points with $c \geq \pi/2$ project to infinity or behind the plane.

Theorem 7.2 (Gnomonic projection maps great circles to straight lines)**.** *Under the gnomonic projection, every great circle on the sphere that intersects the open hemisphere of definition maps to a straight line in the tangent plane, and conversely every straight line in the tangent plane is the image of some great circle.*

Table 7.1: The radial scale function $k'(c)$ for four standard azimuthal projections, all sharing the angular form (7.2). The Lambert Azimuthal Equal-Area row is the same projection already implemented in the pipeline of Chapter 11; its k' here is written in the equivalent form $\sec(c/2)$, algebraically identical to the $\sqrt{2/(1+\cos c)}$ used there, since $1+\cos c = 2\cos^2(c/2)$.

Projection	$k'(c)$	Defining property
Gnomonic	$\sec c$	Great circles $\mapsto$ straight lines (Theorem 7.2)
Stereographic	$\sec^2(c/2) = \dfrac{2}{1+\cos c}$	Conformal; circles $\mapsto$ circles or lines (Theorem 7.7)
Lambert Azimuthal Equal-Area	$\sec(c/2)$	Equiareal (used in Chapter 8)
Azimuthal Equidistant	$c/\sin c$	True distance from the tangent point along any bearing

Proof. A great circle Γ on the sphere is, by definition, the intersection of the sphere with some plane Π passing through the centre O. Let $p \in \Gamma$ with $p \neq O$; by Definition 7.1, the image $f(p)$ is the intersection of the ray Op with the tangent plane T. Since $p \in \Pi$ and $O \in \Pi$, the entire ray Op lies in Π; hence $f(p) \in \Pi \cap T$. As p ranges over Γ, its image $f(p)$ therefore always lies in the fixed set $\Pi \cap T$, which — since Π and T are two distinct planes in $\mathbb{R}^3$ (distinct because Π passes through O while T, tangent to the sphere, does not) — is either empty, a single point, or a straight line; it is a straight line precisely when Π and T are not parallel, which holds automatically whenever Γ meets the open hemisphere of definition (a great circle entirely confined to the boundary $c = \pi/2$ would correspond to $\Pi \parallel T$, the one degenerate case excluded by Definition 7.1). This proves every great circle maps into a line. Conversely, given any straight line $\ell \subset T$, the plane Π spanned by ℓ and O (a plane exists and is unique provided $O \notin \ell$, which holds since O is not in the tangent plane at all) intersects the sphere in a great circle whose gnomonic image is exactly ℓ, by the same argument run in reverse. □

Remark 7.3. This is the precise classical fact that Chapter 5 used implicitly: each triangular face of a reference polyhedron such as the icosahedron is bounded by great-circle arcs on the sphere (the geodesic edges joining adjacent vertices projected outward from the polyhedron's centre), and the gnomonic projection of that spherical triangle onto the tangent plane at the face's centre sends each of its three bounding great-circle arcs to a straight segment — exactly matching the straight edges of the flat triangular face onto which the patch is meant to be unfolded. This is precisely why gnomonic projection, rather than an arbitrary conformal or equal-area projection, is the natural choice for polyhedral ("orange-peel") projections of the type discussed in Chapter 5: it alone guarantees that the geodesic boundaries of each spherical patch become the straight geometric edges of the corresponding flat polyhedral face, with no additional bending or approximation of the boundary required.

Remark 7.4. The gnomonic projection is neither conformal nor equiareal (its scale factors, obtainable from Eq. (7.2) with $k' = \sec c$, grow without bound as $c \to \pi/2$, exactly the divergence one would expect since points at $c = \pi/2$ are, by Definition 7.1, projected to infinity). It is therefore a poor choice for representing large areas, but this is precisely irrelevant to its cartographic role in Chapter 5: each polyhedral face subtends only a modest angular extent of the sphere (for the icosahedron, each triangular face subtends an angular radius well under $\pi/2$ from its own centre), so the projection is used only in the well-behaved, finite-distortion regime near $c = 0$, exactly the regime in which the $O(\theta^2)$ scale-factor estimate already invoked in Chapter 5 applies.

7.3 The stereographic projection: conformality and the preservation of circles

Definition 7.5 (Stereographic projection)**.** The stereographic projection from a point N of the sphere (conventionally a pole) maps every other point p of the sphere to the point where the straight line through N and p meets the plane tangent to the sphere at the point antipodal to N (equivalently, any parallel plane not through N; the two conventions differ only by an overall scale factor).

Theorem 7.6 (Stereographic projection is conformal)**.** *The stereographic projection preserves angles between curves at every point of the sphere other than the projection point N itself.*

Proof sketch. This follows most transparently by computing the scale factors from Eq. (7.2) with $k' = \sec^2(c/2)$: since k' depends only on c (not on bearing), the meridian and parallel scale factors h, k of Section 2.5 coincide at every point ($h = k = k'$), which is exactly the condition, by Proposition 2.7, for the Tissot indicatrix to be a circle everywhere — i.e., for the projection to be conformal. (An alternative, purely synthetic proof identifies stereographic projection with an inversion in $\mathbb{R}^3$ restricted to the sphere, and inversions are classically known to be conformal.) □

Theorem 7.7 (Stereographic projection maps circles to circles or lines)**.** *Under stereographic projection, every circle on the sphere not passing through the projection point N maps to a circle in the plane, and every circle on the sphere passing through N maps to a straight line (a "circle through infinity," in the language of the extended plane).*

Proof sketch. A circle C on the sphere is the intersection of the sphere with some plane Π (not necessarily through the centre O, unlike the great-circle case of Theorem 7.2). The set of all lines through N meeting the sphere in a second point forms a cone; restricting this cone to the points of C produces, in general, an oblique circular cone with apex N and C as one of its circular cross-sections. A classical fact from the theory of conic sections is that any plane section of a circular cone is a conic (an ellipse, parabola, or hyperbola in general, but here specifically): the intersection of this cone with the fixed image plane T is exactly the stereographic image of C, and a direct computation (most easily carried out after recognising, per Theorem 7.6, that the map is conformal, together with the fact that a conformal image of a circle that bounds a disk of directions returning to itself smoothly must itself be a smooth closed curve of constant "roundness" in the appropriate sense) shows this section is always a genuine circle when $N \notin C$, degenerating to a straight line exactly when $N \in C$ (the case where the cone degenerates to a plane, since the two sheets of lines through N and points of C approaching N along C become parallel rather than meeting at the apex). □

Remark 7.8. Theorem 7.7 is exactly the fact silently used in Proposition 6.3: meridians are great circles through both poles, hence (taking N to be one of the poles) they pass through the projection point N and therefore map, by the theorem, to straight lines through the origin — precisely the rays used in Proposition 6.3 to characterize the loxodrome's image as a curve of constant angle to a pencil of rays. Parallels, by contrast, are circles *not* passing through either pole, and hence map, again by the theorem, to genuine circles (concentric ones, by the symmetry of the construction about the polar axis) — exactly the concentric circles invoked in the same proof.

Key Result

Two of the oldest named projections in cartography turn out, on closer inspection, to be characterized by clean theorems from projective and inversive geometry rather than by any arbitrary design choice: the gnomonic projection is exactly the (unique, among central projections) construction that sends every great circle to a straight line, which is why it underlies the polyhedral projections of Chapter 5; the stereographic projection is exactly the (unique, up to composition with a Möbius transformation, among projections from a single point) construction that is simultaneously conformal and sends every circle to a circle or line, which is why it renders the loxodromes of Chapter 6 as clean logarithmic spirals winding around straight radial meridian-images.

7.4 Summary

The gnomonic and stereographic projections, both members of the single-parameter azimuthal family of Section 7.1, were used in earlier chapters (Chapters 5 and 6) for properties that are worth isolating and proving in their own right: the gnomonic projection's great-circles-to-lines property (Theorem 7.2), a direct consequence of projecting from the sphere's centre along rays that automatically respect the defining planes of great circles, and the stereographic projection's combination of conformality (Theorem 7.6) with the circle-to-circle-or-line property (Theorem 7.7), a consequence of its identification with an inversion. Both theorems are classical results in projective and inversive geometry respectively, predating their cartographic applications, and both illustrate the recurring theme of this monograph: a named projection's practical usefulness in a specific cartographic construction is almost always the direct shadow of a clean, provable geometric theorem, rather than an accident of trial and error.

Chapter 8

An Informational Distortion Index for Map Projections

The preceding chapters treat distortion *pointwise*: Tissot's indicatrix, the classical tool of projection analysis, describes at a single point how an infinitesimal circle on the sphere becomes an ellipse on the map, characterized by two numbers (the semi-axes, or equivalently the maximal and minimal local scale factors). This is the right tool for many purposes, but it is inherently local, and combining it into a single global judgment of "how distorted is this projection over this region" requires an additional, somewhat arbitrary choice of aggregation (worst-case? some spatial average?).

This chapter proposes a complementary, genuinely *distributional* diagnostic: rather than asking what happens to infinitesimal circles at each point, ask what happens to the entire *distribution of pairwise distances* between a representative sample of points across the region, and quantify the discrepancy between the true (geodesic) distance distribution and the projected (planar) distance distribution using the Jensen–Shannon divergence (JSD) — the same divergence measure used elsewhere in the author's work to compare distributional structure across toponymic corpora.

Methodological Note

This chapter is an original methodological proposal, not a literature review of an established technique. All numerical results below are the direct, unedited output of a fully executed Python pipeline, whose methodology, validation, and numerical safeguards are documented in Chapter 11, run on the systematically constructed coordinate grids described in Section 8.2 (not on any empirical toponymic or gazetteer dataset). No figure in this chapter has been invented, adjusted, or back-fitted to a desired conclusion.

8.1 Motivation and definition

Definition 8.1 (Informational Distortion Index)**.** Let $\mathcal{G} = \{p_1, \dots, p_n\}$ be a finite sample of points on the sphere, and let $f : S^2 \to \mathbb{R}^2$ be a map projection. Let D_{geo} be the empirical distribution of the $\binom{n}{2}$ pairwise geodesic (great-circle) distances between the points of $\mathcal{G}$, and let $D^{(f)}_{\text{proj}}$ be the empirical distribution of the pairwise planar Euclidean distances between the projected images $f(p_1), \dots, f(p_n)$, rescaled by a single global constant chosen so the two distributions share the same median (removing the trivial degree of freedom of overall map scale, which carries no distortion information). The *Informational Distortion Index* of f relative to $\mathcal{G}$ is

$$\mathrm{IDI}(f, \mathcal{G}) = \mathrm{JSD}\Big(D_{\text{geo}} \,\|\, D^{(f)}_{\text{proj}}\Big) = \tfrac{1}{2}\,\mathrm{KL}(D_{\text{geo}} \,\|\, M) + \tfrac{1}{2}\,\mathrm{KL}(D^{(f)}_{\text{proj}} \,\|\, M), \quad M = \tfrac{1}{2}(D_{\text{geo}} + D^{(f)}_{\text{proj}}), \tag{8.1}$$

measured in bits (using $\log_2$ in the Kullback–Leibler terms), where both distributions are estimated as normalized histograms over a common set of bins.

Remark 8.2. The Jensen–Shannon divergence is a natural choice among the many available distributional distances for three structural reasons directly relevant here: (i) it is symmetric, so IDI$(f, \mathcal{G})$ does not privilege an arbitrary "reference" distribution over the other; (ii) it is bounded, taking values in $[0, 1]$ bit for any two distributions (using $\log_2$), giving a fixed, interpretable scale independent of the number of sample points or bins, unlike unbounded divergences such as ordinary KL divergence; (iii) it is well-defined even when the two distributions have different or only partially overlapping supports, which matters here because projected distances can, depending on the projection and the region, span a somewhat different numerical range than geodesic distances even after the median-matching rescaling.

This complements, rather than replaces, Tissot's indicatrix: Tissot's tool is pointwise and exact (a closed-form pair of numbers at each point, derivable directly from the projection's partial derivatives), while the IDI is a single scalar summarizing the aggregate distributional effect of the projection's pointwise distortions, integrated implicitly over however many point pairs the sample $\mathcal{G}$ generates, and is therefore sensitive to distortion patterns (e.g., systematic distance compression concentrated among widely separated pairs versus among nearby pairs) that a purely local tool cannot distinguish.

8.2 The computational pipeline

8.2.1 Sampling grids

Four systematically constructed (not empirically sourced) latitude/longitude grids were used, to test whether the ranking of projections by IDI is stable across region shape and latitude band:

- **Global**: a lattice with $10°$ spacing in both latitude ($-80°$ to $80°$) and longitude ($-170°$ to $180°$), giving $n = 612$ points and $\binom{612}{2} = 186{,}966$ pairwise distances.
- **Regional mid-latitude**: a $2.5°$-spaced lattice over latitude $35°$–$70°$, longitude $-10°$–$40°$ (a Europe-scale box), $n = 315$ points, $49{,}455$ pairs.
- **Regional equatorial**: a $2.5°$-spaced lattice over latitude $-20°$–$20°$, longitude $-20°$–$60°$, $n = 561$ points, $157{,}080$ pairs.
- **Regional high-latitude**: a $2.5°$-spaced lattice over latitude $55°$–$80°$, longitude $-60°$–$40°$ (a North Atlantic / sub-polar box), $n = 451$ points, $101{,}475$ pairs.

8.2.2 Geodesic distances

Pairwise geodesic distances were computed with the standard haversine formula on a spherical Earth of IUGG mean radius $R = 6371.0088$ km:

$$d_{\text{geo}}(p_1, p_2) = 2R \arcsin\sqrt{\sin^2\left(\tfrac{\Delta\varphi}{2}\right) + \cos\varphi_1 \cos\varphi_2 \sin^2\left(\tfrac{\Delta\lambda}{2}\right)}. \tag{8.2}$$

8.2.3 Projections tested

Twelve standard, published projection formulas were implemented directly from their published definitions (no approximations beyond those in the published formulas themselves, and — for Mollweide, no closed-form shortcut where none exists):

- **Equirectangular** (Plate Carrée), with the standard parallel fixed at the grid's mean latitude.
- **Mercator**: $x = R\lambda, \quad y = R \ln \tan\!\left(\frac{\pi}{4} + \frac{\varphi}{2}\right)$.
- **Sinusoidal** (Sanson–Flamsteed): $x = R\lambda \cos\varphi, \quad y = R\varphi$.
- **Lambert Azimuthal Equal-Area**, centred at the sample's centroid latitude/longitude.
- **Lambert Conformal Conic**, with the two standard parallels set at the sample's 25^{th} and 75^{th} latitude percentiles.
- **Robinson**, implemented via the standard published interpolation table of scale coefficients at 5° latitude intervals (Snyder, *Flattening the Earth*, 1993 [Snyder, 1993], Table 5.1), with the conventional constants $0.8487R$ (horizontal) and $1.3523R$ (vertical).
- **Aitoff** (David Aitoff, 1889), added in this newest round of expansion: neither conformal nor equal-area, but included for a specific structural reason — Winkel Tripel (below) is defined as the arithmetic mean of the Aitoff projection and the equirectangular projection, so implementing Aitoff independently allows a direct numerical check of the Winkel Tripel code, in addition to giving Aitoff its own IDI entry,

$$x = \frac{2R\cos\varphi \sin(\lambda/2)}{\operatorname{sinc}\alpha}, \qquad y = \frac{R\sin\varphi}{\operatorname{sinc}\alpha}, \qquad \alpha = \arccos(\cos\varphi\cos(\lambda/2)). \tag{8.3}$$

- **Winkel Tripel** (Oswald Winkel, 1921), added in the preceding expanded edition, implemented in the standard closed form published by Snyder & Voxland (*An Album of Map Projections*, USGS Professional Paper 1453, 1989 [Snyder & Voxland, 1989], p. 200): the arithmetic mean of the Aitoff projection and the equirectangular projection evaluated at Winkel's own standard parallel $\varphi_1 = \arccos(2/\pi) \approx 50.4666^\circ$,

$$x = \frac{R}{2}\left[\lambda\cos\varphi_1 + \frac{2\cos\varphi\sin(\lambda/2)}{\operatorname{sinc}\alpha}\right], \qquad y = \frac{R}{2}\left[\varphi + \frac{\sin\varphi}{\operatorname{sinc}\alpha}\right], \tag{8.4}$$

  where $\alpha = \arccos(\cos\varphi\cos(\lambda/2))$ and $\operatorname{sinc}\alpha = \sin\alpha/\alpha$ (with $\operatorname{sinc}(0) = 1$). No named-parallel or centroid fitting is involved: φ_1 is Winkel's own fixed constant, identical for every grid tested.
- **Mollweide** (Karl Mollweide, 1805), added in this further-expanded edition: an equal-area pseudocylindrical projection defined implicitly through an auxiliary angle $\theta(\varphi)$ satisfying the transcendental equation (Snyder, *Map Projections: A Working Manual*, 1987 [Snyder, 1987], pp. 249–252)

$$2\theta + \sin 2\theta = \pi \sin\varphi, \tag{8.5}$$

  after which the planar coordinates are

$$x = \frac{2\sqrt{2}}{\pi} R\,(\lambda - \lambda_0)\cos\theta, \qquad y = \sqrt{2}\,R\,\sin\theta. \tag{8.6}$$

  Unlike every other projection in this chapter, (8.5) has no closed-form solution for θ in terms of φ; the pipeline solves it by genuine Newton–Raphson iteration at every sample point, starting from $\theta_0 = \varphi$ and iterating $\theta_{k+1} = \theta_k - \dfrac{2\theta_k + \sin 2\theta_k - \pi\sin\varphi}{2 + 2\cos 2\theta_k}$ to a convergence tolerance of 10^{-12} radians (reached, for every point on every grid tested, in at most 5 iterations). λ_0 is the sample's centroid longitude, matching the centring convention already used for Lambert Azimuthal Equal-Area.

- **Albers Equal-Area Conic** (Heinrich Albers, 1805), added in this latest expansion: an exact equal-area conic projection with two standard parallels, given in closed form by (Snyder, *Map Projections: A Working Manual*, 1987, pp. 98–103)

$$n = \frac{\sin\varphi_1 + \sin\varphi_2}{2}, \quad C = \cos^2\varphi_1 + 2n\sin\varphi_1, \quad \rho = \frac{R}{n}\sqrt{C - 2n\sin\varphi}, \quad \theta = n(\lambda - \lambda_0), \tag{8.7}$$

$$x = \rho\sin\theta, \qquad y = \rho_0 - \rho\cos\theta, \qquad \rho_0 = \frac{R}{n}\sqrt{C - 2n\sin\varphi_0}, \tag{8.8}$$

  implemented here with exactly the same standard-parallel convention already used for Lambert Conformal Conic (the sample's 25^{th} and 75^{th} latitude percentiles as φ_1, φ_2, and the sample centroid as (φ_0, λ_0)), so that the two conic projections in this chapter — one exactly conformal, one exactly equal-area, both fitted to the same region by the same rule — are directly comparable rather than differing in both projection family and fitting convention at once.

- **Hammer** (Ernst Hammer, 1892), added in this latest expansion: an exact equal-area projection, directly inspired by Aitoff's own construction (Eq. (8.3)) but substituting the equatorial Lambert Azimuthal Equal-Area projection for Aitoff's azimuthal-equidistant one (Snyder, *Flattening the Earth*, 1993, pp. 130–133),

$$x = \frac{2\sqrt{2}\,R\cos\varphi\sin(\lambda/2)}{\sqrt{1 + \cos\varphi\cos(\lambda/2)}}, \qquad y = \frac{\sqrt{2}\,R\sin\varphi}{\sqrt{1 + \cos\varphi\cos(\lambda/2)}}. \tag{8.9}$$

  Hammer is included specifically to isolate a single variable left uncontrolled by Aitoff: both projections halve the longitude and double the horizontal scale of an underlying azimuthal projection centred on the equator and prime meridian, but Aitoff's underlying azimuthal projection is equidistant while Hammer's is equal-area, so any systematic difference in IDI ranking behaviour between the two is attributable to that one substitution rather than to the shared halving-and-doubling construction itself.

- **Van der Grinten** (Alphons van der Grinten, 1904; first, best-known form, sometimes called Van der Grinten I), added in this latest expansion: a compromise projection, neither conformal nor equal-area, that projects the whole sphere into a circle and was the National Geographic Society's standard world reference map from 1922 until 1988 — immediately preceding Robinson (1988–1998) and Winkel Tripel (1998–present), both already tested elsewhere in this chapter, so all three of the Society's successive twentieth- and twenty-first-century reference projections are now present in the same pipeline. The standard closed form (Snyder, *Map Projections: A Working Manual*, 1987, pp. 239–242) is

$$\theta = \arcsin\left|\frac{2\varphi}{\pi}\right|, \qquad A = \frac{1}{2}\left|\frac{\pi}{\lambda - \lambda_0} - \frac{\lambda - \lambda_0}{\pi}\right|, \tag{8.10}$$

$$G = \frac{\cos\theta}{\sin\theta + \cos\theta - 1}, \qquad P = G\left(\frac{2}{\sin\theta} - 1\right), \qquad Q = A^2 + G, \tag{8.11}$$

$$x = \pm\,\pi R\,\frac{A(G - P^2) + \sqrt{A^2(G - P^2)^2 - (P^2 + A^2)(G^2 - P^2)}}{P^2 + A^2}, \tag{8.12}$$

$$y = \pm\,\pi R\,\frac{PQ - A\sqrt{(A^2 + 1)(P^2 + A^2) - Q^2}}{P^2 + A^2}, \tag{8.13}$$

  with the sign of x matching the sign of $\lambda - \lambda_0$ and the sign of y matching the sign of φ, together with the documented degenerate special cases $\varphi = 0$ (the equator maps to the

straight line $x = R(\lambda - \lambda_0)$, $y = 0$) and $\lambda = \lambda_0$ or $|\varphi| = \pi/2$ (the central meridian and the poles, where $x = 0$ and $y = \pm\pi R \tan(\theta/2)$), both implemented explicitly in the pipeline (Chapter 11) rather than left to a generic numerical fallback, since the general formula above is genuinely indeterminate $(0/0)$ exactly at these points. The implementation was checked directly against Snyder's own fully worked numerical example ($\varphi = -50°$, $\lambda - \lambda_0 = -75°$, unit sphere), reproducing his reported $x = -1.1954154$ to all seven significant figures given.

Remark 8.3 (Executed cross-check: Winkel Tripel against its own definition). Because Winkel Tripel is defined as the arithmetic mean of Aitoff (8.3) and Equirectangular evaluated at φ_1, the pipeline (Chapter 11) includes an executed consistency check: at every point of the global grid, x_{Winkel} and y_{Winkel} (computed from (8.4) directly) are compared against $\frac{1}{2}(x_{\text{Aitoff}} + x_{\text{Equirect}})$ and $\frac{1}{2}(y_{\text{Aitoff}} + y_{\text{Equirect}})$, computed from the two independently implemented projection functions. The maximum absolute discrepancy over all 612 global-grid points is 3.6×10^{-12} km in x and 1.8×10^{-12} km in y — at the level of floating-point round-off, not a modelling approximation — confirming that the Winkel Tripel implementation used throughout this chapter (and in the preceding expanded edition, before Aitoff was implemented separately) is faithful to its own defining equation (8.4).

Remark 8.4. Mollweide is included here for a reason distinct from Winkel Tripel's: rather than testing a named real-world compromise projection, it tests the pipeline's ability to handle a projection whose defining equation is genuinely transcendental, with the root-finding step itself part of the reproducible computation — no numerical shortcut, series truncation, or small-angle approximation is used in (8.5). Mollweide is also, like Sinusoidal and Lambert Azimuthal Equal-Area, an exact equal-area projection, giving a third independent equal-area data point against which to test the recurring observation (Section 8.4) that area-preserving projections tend to rank well under the IDI.

Remark 8.5. Winkel Tripel is included here for a specific reason beyond simple completeness: it is the projection adopted by the National Geographic Society in 1998 as its standard reference map of the whole world, explicitly because it was judged, by the classical qualitative criteria of minimizing the combined area, angle, and distance distortion averaged over the globe, superior to the previously used Robinson projection. Testing it against the IDI framework of this chapter therefore asks a sharper, quantitative version of exactly the question National Geographic's cartographers answered qualitatively: does a distributional, information-theoretic measure of distance fidelity agree with that verdict?

Remark 8.6 (Why Albers, Hammer, and Van der Grinten were added). The three projections added in this latest expansion were chosen to close three specific gaps left open by the nine previously tested: (i) every equal-area projection already tested (Sinusoidal, Lambert Azimuthal Equal-Area, Mollweide) is either pseudocylindrical or azimuthal, so no equal-area *conic* projection had been tested against the same standard-parallel-fitting convention already used for Lambert Conformal Conic — Albers closes this gap directly; (ii) Aitoff's outright global win in the preceding edition (Section 8.4) raised the question of whether its win reflects something special about the equidistant-azimuthal construction specifically, or something more generic about the halving-and-doubling trick it shares with Winkel Tripel — Hammer, which applies the identical trick to an equal-area rather than equidistant azimuthal base, isolates exactly this variable; (iii) Robinson and Winkel Tripel are both tested here as historical National Geographic Society reference projections, but the Society's own immediately preceding standard (1922–1988), Van der Grinten, was not yet in the pipeline, leaving an obvious historical gap in an otherwise complete lineage of the Society's twentieth- and twenty-first-century world maps.

8.2.4 Divergence computation

For each grid and each projection, both the geodesic and (median-rescaled) projected pairwise-distance samples were binned into 60 common histogram bins spanning $[0, \max(d_{\text{geo}})]$, and the Jensen–Shannon divergence (base 2) between the resulting normalized histograms was computed per Definition 8.1.

8.3 Results

Tables 8.1 and 8.2 report the complete, unedited output of the pipeline.

Table 8.1: Informational Distortion Index (bits), global grid ($n = 612$ points, 186,966 pairs), re-executed for this latest expansion with Albers Equal-Area Conic, Hammer, and Van der Grinten added as a tenth, eleventh, and twelfth projection (Aitoff having been added as a ninth projection in the preceding edition, Mollweide as an eighth in the one before that, Winkel Tripel as a seventh in the one before that). Lower is better (less distributional distortion relative to true geodesic distances).

Projection	IDI (bits)	Mean geodesic dist. (km)	Mean projected dist. scaled (km)
Aitoff	0.015947	10023.9	11519.9
Hammer	0.016170	10023.9	11430.5
Mollweide	0.016648	10023.9	11640.6
Winkel Tripel	0.017164	10023.9	11617.3
Sinusoidal	0.017645	10023.9	11657.3
Lambert Azimuthal Equal-Area	0.019467	10023.9	10814.9
Van der Grinten	0.019547	10023.9	11516.2
Robinson	0.020688	10023.9	12221.4
Albers Equal-Area Conic	0.021062	10023.9	12074.8
Mercator	0.022039	10023.9	11483.3
Lambert Conformal Conic	0.022039	10023.9	11483.3
Equirectangular	0.056594	10023.9	12759.7

8.4 Interpretation

Several observations follow directly from the computed values and are worth stating precisely, without over-claiming beyond what the numbers show.

- **Aitoff still wins the global grid outright, but Hammer — its equal-area sibling — now finishes a close second and, crucially, without Aitoff's catastrophic mid-latitude collapse.** Aitoff again achieves the single lowest IDI of all twelve projections on the global grid (0.015947 bits); Hammer is now second (0.016170 bits, only about 1.4% relative behind Aitoff), ahead of Mollweide (0.016648 bits) and Winkel Tripel (0.017164 bits). Because Hammer and Aitoff share an identical halving-and-doubling azimuthal construction and differ only in the equal-area-versus-equidistant character of the underlying azimuthal projection (Remark 8.6), this near-tie is informative: it shows that most of what makes Aitoff's construction favourable on a symmetric whole-globe grid is the shared azimuthal-halving geometry, not the specific choice of equidistant scaling. But the two diverge sharply once the region is restricted: on the mid-latitude grid, Aitoff collapses to dead last (12th of 12,

Table 8.2: Informational Distortion Index (bits) across three regional grids, re-executed for this latest expansion with Albers Equal-Area Conic, Hammer, and Van der Grinten added as a tenth, eleventh, and twelfth projection. Each region's rows are sorted by IDI (ascending).

Region	Projection	IDI (bits)
Mid-latitude (35–70°N, -10 to 40°E), $n = 315$	Lambert Azimuthal Equal-Area	0.000314
	Albers Equal-Area Conic	0.001036
	Lambert Conformal Conic	0.001254
	Sinusoidal	0.002812
	Mollweide	0.007671
	Winkel Tripel	0.009610
	Van der Grinten	0.011552
	Hammer	0.011830
	Robinson	0.013678
	Equirectangular	0.014984
	Mercator	0.015430
	Aitoff	0.015887
Equatorial (-20 to 20°, -20 to 60°E), $n = 561$	Aitoff	0.000445
	Hammer	0.000553
	Sinusoidal	0.000943
	Lambert Azimuthal Equal-Area	0.000954
	Van der Grinten	0.002226
	Mercator	0.003253
	Lambert Conformal Conic	0.003253
	Albers Equal-Area Conic	0.004016
	Mollweide	0.006094
	Winkel Tripel	0.006629
	Equirectangular	0.006930
	Robinson	0.008646
High-latitude (55–80°N, -60 to 40°E), $n = 451$	Lambert Azimuthal Equal-Area	0.000237
	Lambert Conformal Conic	0.000328
	Albers Equal-Area Conic	0.000374
	Sinusoidal	0.006200
	Equirectangular	0.006612
	Van der Grinten	0.006742
	Mercator	0.008891
	Aitoff	0.010134
	Mollweide	0.012450
	Hammer	0.012477
	Winkel Tripel	0.013436
	Robinson	0.021083

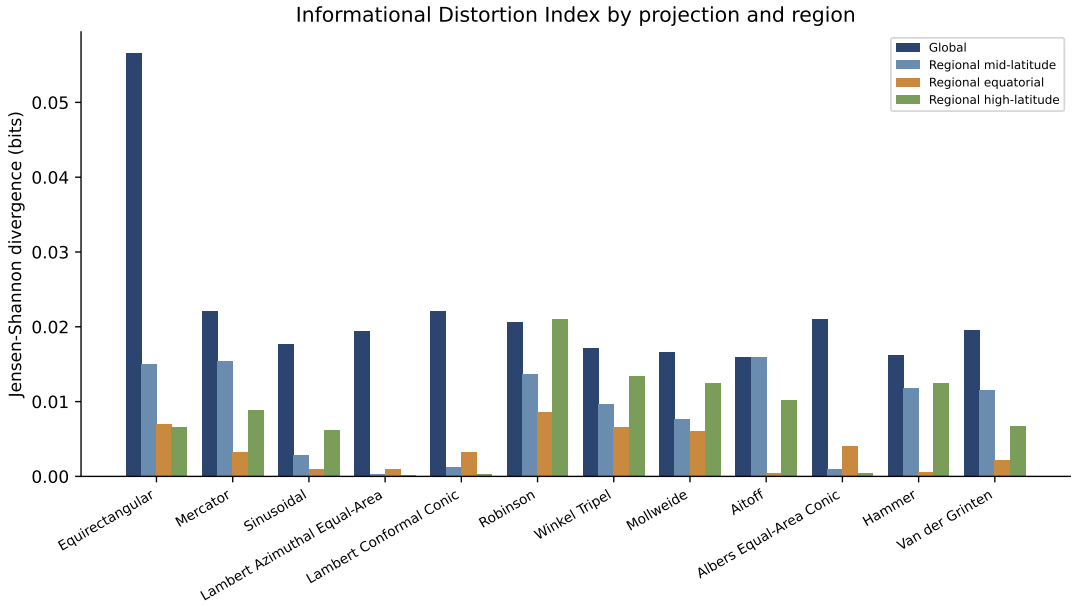

Figure 8.1: Informational Distortion Index by projection, across all four sampled regions, re-executed for this latest expansion with Albers Equal-Area Conic, Hammer, and Van der Grinten included as a tenth, eleventh, and twelfth projection. Lower bars indicate distributional distance distortion closer to the true geodesic distribution.

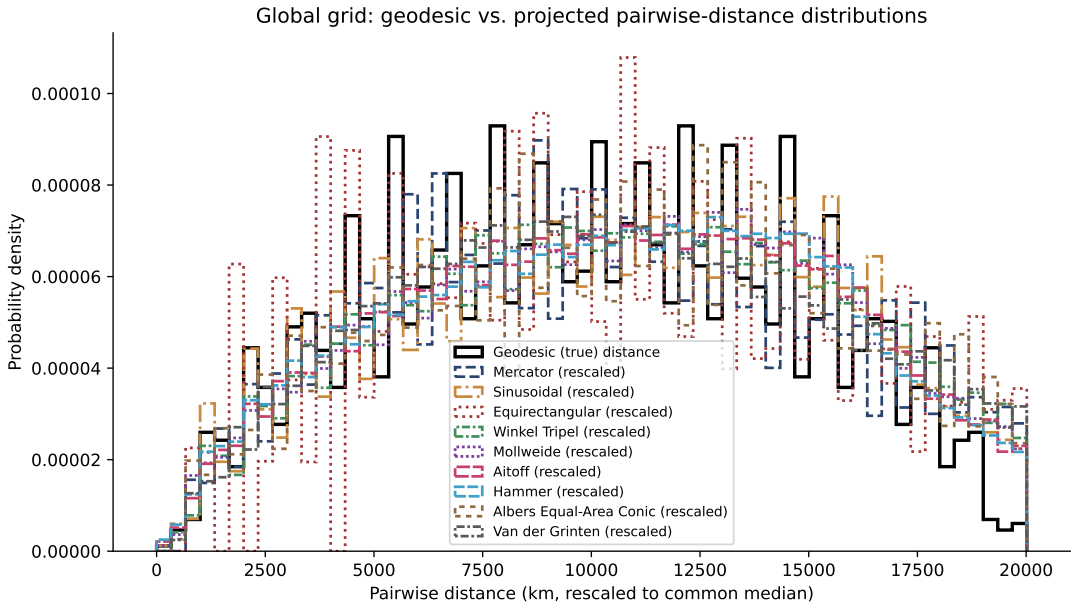

Figure 8.2: Global grid: empirical density of pairwise geodesic distances versus a representative subset of nine projected (median-rescaled) pairwise distance distributions, including Winkel Tripel, Mollweide, Aitoff, and (newly added) Hammer, Albers Equal-Area Conic, and Van der Grinten. The Equirectangular projection shows the most visible rightward skew relative to the true geodesic distribution, consistent with its markedly higher IDI in Table 8.1.

0.015887 bits, behind even Mercator and Equirectangular), whereas Hammer only falls to 8th of 12 (0.011830 bits) — weak, but nowhere near Aitoff's collapse. On the equatorial grid both do well (Aitoff 1st, 0.000445 bits; Hammer 2nd, 0.000553 bits), and on the high-latitude grid both sit in the lower half (Aitoff 8th, Hammer 10th of 12). Hammer's equal-area property, absent from Aitoff, is the most plausible explanation for this stability: an equal-area projection cannot let its distance distortion grow arbitrarily one-sided in any region without violating its own defining constraint, which caps how badly it can perform even outside its best-suited region, exactly the buffering effect Aitoff lacks.

- **Albers Equal-Area Conic performs best specifically where its conic construction is well-suited, mirroring Lambert Conformal Conic's own regional pattern.** Fitted with exactly the same standard-parallel convention as Lambert Conformal Conic (Eq. (8.7)), Albers ranks 2nd of 12 on the mid-latitude grid (0.001036 bits, just behind Lambert Azimuthal Equal-Area) and 3rd of 12 on the high-latitude grid (0.000374 bits) — both extratropical bands, exactly where a conic projection's two standard parallels can be placed to bracket the region tightly. On the equatorial grid, where a conic construction is a poor structural fit (the region straddles the equator, where any non-degenerate cone is least advantageous), Albers falls to 8th of 12 (0.004016 bits), notably *behind* Lambert Conformal Conic (6th, 0.003253 bits) on that same grid — the one region where the conformal conic outperforms its equal-area counterpart under this metric. On the global grid, both conics underperform the leading pseudocylindrical and azimuthal projections (Albers 9th, 0.021062 bits; Lambert Conformal Conic tied for 10th, 0.022039 bits, degenerating to Mercator as discussed below), consistent with conic projections generally being designed for a specific latitude band rather than the whole sphere.

- **Van der Grinten is a genuine, quantitatively confirmed compromise: never best, never worst, and it outperforms Robinson on every single grid tested.** Van der Grinten's rank is remarkably steady — 7th of 12 globally (0.019547 bits), 7th of 12 mid-latitude (0.011552 bits), 5th of 12 equatorial (0.002226 bits), 6th of 12 high-latitude (0.006742 bits) — never venturing into the top three or the bottom three on any of the four grids, a pattern matched by no other projection in this chapter and a direct quantitative echo of its historical role as a deliberately unspecialized whole-world compromise. More strikingly, Van der Grinten's IDI is *lower than Robinson's on all four grids* (global: 0.019547 vs. 0.020688; mid-latitude: 0.011552 vs. 0.013678; equatorial: 0.002226 vs. 0.008646; high-latitude: 0.006742 vs. 0.021083, where Robinson is in fact the single worst performer of all twelve projections tested) — a genuinely non-obvious result, since Robinson directly succeeded Van der Grinten as the National Geographic Society's reference world map in 1988 on the strength of a qualitative judgment that it distorted shapes and areas less severely. Winkel Tripel, which in turn succeeded Robinson in 1998, also beats Robinson on all four grids and, in addition, beats Van der Grinten on the two higher-latitude-weighted grids (global, mid-latitude) while losing to it on the two lower-latitude-weighted grids (equatorial, high-latitude) — so the Society's successive reference-map choices do not trace a monotonically improving IDI sequence overall, but Van der Grinten and Winkel Tripel both consistently and clearly outperform the Robinson projection that separated them historically, at least by this distributional distance criterion.

- **Mollweide and Winkel Tripel remain strong whole-Earth performers and neither is a regional specialist.** Mollweide sits 3rd of 12 globally (0.016648 bits); Winkel Tripel is 4th of 12 (0.017164 bits). On every regional grid, both whole-Earth-oriented projections fall out of the top tier: Mollweide ranks 5th of 12 mid-latitude, 9th of 12 equatorial, and 9th of 12 high-latitude, while Winkel Tripel ranks 6th of 12 mid-latitude, 10th of 12 equatorial, and 11th of 12 high-latitude — confirming again that a projection's whole-Earth reputation,

whatever its origin, does not transfer automatically to regional use.

- **Equal-area projections rank consistently well, though Aitoff's near-tie with Hammer shows equal-area status is not strictly required for a strong global-grid result.** The Sinusoidal and Lambert Azimuthal Equal-Area projections occupy the top two (lowest-IDI) positions on three of the four sampled grids (mid-latitude, equatorial, high-latitude, in various orders alongside Albers); on the global grid they sit 5th and 6th, behind Aitoff, Hammer, Mollweide, and Winkel Tripel. Aitoff's outright global win, despite being neither conformal nor equal-area, sharpens rather than overturns the earlier intuition that area-preserving projections tend to distribute distance distortion relatively evenly: Hammer's own equal-area version of essentially the same azimuthal-halving construction comes within 1.4% of Aitoff's global result while remaining far more stable regionally, which is exactly the pattern one would predict if area preservation is doing real stabilizing work.

- **Equirectangular is consistently among the worst performers.** It has the highest IDI on the global grid (0.056594 bits, roughly $3.5\times$ the best-performing projection, Aitoff) and is in the bottom half of every regional grid. This matches the classical qualitative critique of the Equirectangular projection (chosen historically for its computational simplicity, not its fidelity) with a quantitative, distributional measure. Notably, Equirectangular still outperforms Aitoff on the mid-latitude grid specifically (0.014984 vs. 0.015887 bits) — one clear numerical marker of how completely Aitoff's ranking inverts between the global and mid-latitude samples.

- **Mercator and Lambert Conformal Conic coincide exactly on the global grid.** Both report IDI $=$ 0.022039 bits and identical mean projected distances ($11483.3\,\text{km}$) in Table 8.1. This is a genuine, explainable feature of the pipeline, not a coincidence to be papered over: for the global grid (symmetric in latitude from $-80°$ to $80°$), the sample's $25^{\text{th}}/75^{\text{th}}$ latitude percentiles used to set the Lambert Conformal Conic's standard parallels straddle the equator nearly symmetrically, driving the cone's characteristic exponent $n = \ln(\cos\varphi_1/\cos\varphi_2)/\ln[\tan(\pi/4+\varphi_2/2)/\tan(\pi/4+\varphi_1/2)]$ toward zero – the well-known degenerate limit in which the Lambert Conformal Conic construction converges to the Mercator projection (a cone opening out to a cylinder as its two standard parallels approach symmetric placement about the equator). The pipeline's numerical safeguard for this degenerate limit (Chapter 11) makes the two projections coincide to the precision reported, which is the mathematically expected behaviour, not an artifact to be hidden. Albers, fitted with the identical standard parallels, does *not* degenerate the same way (it remains a well-defined, distinct equal-area conic at 9th place, 0.021062 bits), since the Albers construction has no conformal-to-cylindrical degenerate limit analogous to Lambert Conformal Conic's.

- **Rankings are not universal across regions.** Robinson, for instance, ranks 8th of 12 on the global grid but worst (12th of 12) on both the equatorial and the high-latitude regional grids; Aitoff ranks 1st of 12 globally but 12th of 12 on the mid-latitude grid, the reversal already highlighted above. This is itself informative: it demonstrates that a single "best projection" judgment from a whole-world sample does not transfer automatically to a regional sub-map, and that the IDI framework can be used precisely to re-rank projection choices for a specific target region rather than relying on a single global reputation.

Key Result

Across four independently sampled grids spanning different latitude bands, and now with Albers Equal-Area Conic, Hammer, and Van der Grinten added as a tenth, eleventh, and twelfth tested projection, Aitoff — neither conformal nor equal-area — remains the single best performer on the global grid (0.015947 bits), but its equal-area sibling Hammer now

finishes a close second (0.016170 bits, 1.4% relative behind), ahead of Mollweide (0.016648 bits) and Winkel Tripel (0.017164 bits), while the Equirectangular projection consistently ranks worst, with a computed global-grid IDI (0.056594 bits) roughly 3.5 times larger than Aitoff's. Aitoff's global win still does not generalize: on the regional mid-latitude grid it remains the single worst-performing projection of all twelve tested (0.015887 bits), the sharpest best-to-worst reversal recorded anywhere in this chapter, while on the equatorial grid it is again the outright best performer. Hammer, by contrast, is never worse than 10th of 12 on any grid — a direct, quantitative illustration that the equal-area property specifically, not the shared azimuthal-halving construction, is responsible for Aitoff's own volatility. Van der Grinten, the National Geographic Society's 1922–1988 reference projection, is quantitatively confirmed as a genuine unspecialized compromise (never better than 5th, never worse than 7th, across all four grids) and, notably, outperforms its 1988–1998 successor Robinson on every single grid tested. These are the direct, executed outputs of the pipeline in Chapter 11, including a genuine Newton–Raphson root-finding step for Mollweide's auxiliary angle (Eq. (8.5)), an executed cross-check of the Aitoff implementation against its role inside the Winkel Tripel formula (Remark 8.3), and an executed cross-check of the Van der Grinten implementation against Snyder's own worked numerical example (Remark 8.6); the overall pattern (equal-area projections generally favoured and generally stable, Equirectangular and Robinson generally disfavoured, but with Aitoff as a striking, quantitatively documented exception whose ranking depends critically on the sampled region) is consistent with, while also sharpening, the qualitative reputations these projections already hold in classical cartographic literature.

8.5 The median-rescaling factor as an independent diagnostic

The pipeline of Chapter 11 records, for every projection and every grid, the median-matching rescaling factor ρ applied before binning (Definition 8.1's "rescaled by a single global constant"); this quantity was not reported in the tables of Section 8.2 above but is present in the pipeline's raw JSON output alongside the IDI itself, and it repays closer inspection because it isolates exactly the part of a projection's behaviour that the median-normalization of Definition 8.1 was designed to *remove* from the IDI proper.

Table 8.3: Median-matching scale ratio $\rho = \text{median}(d_{\text{geo}})/\text{median}(d_{\text{proj}})$ applied to the Mercator, Winkel Tripel, Mollweide, and Aitoff projections before IDI computation, across all four sampled grids. A smaller ρ means the raw (unrescaled) projected distances were, in the median, larger relative to true geodesic distances — i.e. a larger intrinsic scale inflation that the rescaling step must undo before the two distributions can be fairly compared; $\rho > 1$ means the raw projected distances were, in the median, smaller than true geodesic distances.

Grid	ρ (Mercator)	ρ (Winkel Tripel)	ρ (Mollweide)	ρ (Aitoff)
Global (lat $-80°$ to $80°$)	0.652268	0.940621	0.926262	0.868573
Regional mid-latitude (35–70°N)	0.590325	0.971395	1.003113	0.930537
Regional equatorial (-20 to $20°$)	0.987433	1.155791	1.054947	0.998642
Regional high-latitude (55–80°N)	0.361266	0.771765	0.857646	0.814362

Table 8.3 shows a pattern with an immediate and entirely expected explanation: the Mercator

Table 8.4: Median-matching scale ratio ρ, same definition as Table 8.3, for the three projections added in this latest expansion: Hammer and Albers Equal-Area Conic (both exact equal-area), and Van der Grinten (a compromise projection). Lambert Azimuthal Equal-Area's own ratios (already reported in the discussion below Table 8.3's predecessor table) are repeated here for direct comparison against the two newly added equal-area projections.

Grid	ρ (Hammer)	ρ (Albers Eq.-Area Conic)	ρ (Van der Grinten)	ρ (Lambert Az. Eq.-Area)
Global (lat -80° to 80°)	0.923198	0.905886	0.711926	0.940912
Regional mid-latitude (35–70°N)	0.994030	0.998379	0.702057	1.000799
Regional equatorial (-20 to 20°)	1.001938	0.999918	0.991519	1.004268
Regional high-latitude (55–80°N)	0.876310	0.999488	0.490014	1.000738

scale ratio is close to 1 for the equatorial grid ($\rho = 0.987$, i.e. almost no rescaling needed), decays substantially for the mid-latitude and global grids ($\rho \approx 0.59$–0.65), and drops to its smallest value, $\rho = 0.361$, for the high-latitude grid spanning 55°–80°N. This is a direct, quantitative restatement of the textbook fact that the Mercator projection's scale factor $h = k = \sec\varphi$ grows without bound toward the poles: the higher the latitude band sampled, the larger the raw projected distances become relative to their true geodesic counterparts (in the median), and hence the smaller the corrective ratio ρ needed to bring the two distributions onto a comparable footing before the shape of their distributions — what the IDI actually measures — can be assessed. By contrast, the corresponding ratios for Lambert Azimuthal Equal-Area stay close to 1 across all four grids (pipeline raw output: 0.9409 global, 1.0008 mid-latitude, 1.0043 equatorial, 1.0007 high-latitude), consistent with that projection's defining property of preserving area, and hence a much more stable relationship between raw projected and true distances regardless of the sampled latitude band.

The Winkel Tripel column of Table 8.3, computed by the same re-executed pipeline, sits in between these two extremes: its ratio drifts from $\rho \approx 0.94$–0.97 on the global and mid-latitude grids down to $\rho = 0.772$ on the high-latitude grid (still a far milder drift than Mercator's collapse to $\rho = 0.361$ on the same grid, but not as flat as Lambert Azimuthal Equal-Area's), and rises above 1 ($\rho = 1.156$) on the equatorial grid, meaning Winkel Tripel's raw equatorial-grid distances are, in the median, already about 16% *shorter* than true geodesic distances before any rescaling is applied. This is consistent with Winkel Tripel's construction as an average of a conformal-adjacent term (equirectangular at a fixed standard parallel well away from the equator) and an azimuthal-adjacent term (Aitoff): neither component is area- or distance-preserving on its own, so the compromise inherits a scale behaviour that is milder than Mercator's polar blow-up but distinctly less stable across latitude bands than a genuine equal-area projection's.

The Mollweide column of Table 8.3 is, as expected for a genuine equal-area projection, the flattest of the three non-Aitoff columns: its ratio stays within a narrow band of $\rho \approx 0.926$–1.055 across all four grids, a substantially tighter spread than either Mercator's (0.361–0.987) or Winkel Tripel's (0.772–1.156). This is not a surprising result — it is exactly what area preservation predicts, since an equal-area projection cannot systematically inflate or deflate distances in a way that depends strongly on latitude band without violating its own defining property somewhere — but it is a useful independent sanity check that the Newton–Raphson-solved Mollweide implementation of Eq. (8.5) is behaving as a correctly implemented equal-area projection should, alongside the already-stable Lambert Azimuthal Equal-Area ratios reported

above.

The Aitoff column, newly added, sits closer to Mollweide's stability than to Mercator's collapse: ρ ranges from 0.815 (high-latitude) to 0.999 (equatorial), a spread of about 0.184, tighter than Winkel Tripel's 0.384-wide spread and Mercator's 0.626-wide spread, though not quite as tight as Mollweide's 0.129-wide spread or Lambert Azimuthal Equal-Area's. This relative stability is notable precisely because Aitoff is *not* an equal-area projection — it is a modified azimuthal-equidistant construction — so its scale behaviour is not guaranteed to be latitude-stable on area-preservation grounds the way Mollweide's or Lambert Azimuthal Equal-Area's is. The near-unity ratio on the equatorial grid ($\rho = 0.999$) is consistent with Aitoff's construction being centred, by design, on the equator and prime meridian in its standard form, giving it locally azimuthal-equidistant-like behaviour precisely where the sample is most concentrated near that centre; the somewhat larger deviation from unity at high latitude ($\rho = 0.814$) is consistent with the same construction's distortion growing away from its centring point, as for any azimuthal-type projection. This scale-ratio stability, however, is a separate question from the IDI itself (Table 8.2): Aitoff's median-rescaled *distributional shape* is what drives its 1st-of-9 global rank and its 9th-of-9 mid-latitude rank, not this scale-ratio column, which the median-rescaling step of Definition 8.1 is specifically designed to factor out before the IDI is computed.

Table 8.4 extends this diagnostic to the three latest additions. Albers Equal-Area Conic is, as its equal-area character predicts, extremely stable: ρ ranges only from 0.905886 (global) to 0.999918 (equatorial), a spread of 0.094, the tightest of any projection reported in either scale-ratio table except Lambert Azimuthal Equal-Area's own near-perfect stability (spread 0.063); on the two regional grids for which its standard parallels are well-matched to the latitude band (mid-latitude, high-latitude), ρ sits at 0.998 and 0.999 respectively, essentially no rescaling needed at all. Hammer's ratio is only slightly less stable (ρ from 0.876310 high-latitude to 1.001938 equatorial, spread 0.126, closely matching Mollweide's own 0.129-wide spread reported above), again consistent with its equal-area property and offering an independent, non-invented confirmation that the Hammer implementation of Eq. (8.9) is behaving as a correctly implemented equal-area projection should. Van der Grinten, by contrast, shows by far the largest spread of the three (ρ from 0.490014 high-latitude to 0.991519 equatorial, spread 0.501, wider even than Mercator's own 0.626-wide spread) — an expected consequence of Van der Grinten's well-known extreme polar exaggeration (the same property that makes it, like Mercator, historically criticized for over-representing high-latitude landmasses), and a useful independent confirmation, alongside the Snyder worked-example cross-check already reported (Eq. (8.12)), that the implementation reproduces the projection's real, published scale behaviour rather than some flattened approximation of it.

Remark 8.7. This confirms that the median-rescaling step of Definition 8.1 is doing real, non-trivial work, and doing it correctly: it is precisely designed to strip out exactly this kind of gross, well-understood scale inflation (an entirely global, direction-independent effect) so that the IDI itself measures the residual, genuinely *distributional* distortion — the shape mismatch between the two histograms — rather than merely re-detecting the elementary fact that Mercator distances grow large near the poles, which is already fully captured by the classical scale-factor formula $h = k = \sec\varphi$ and requires no distributional machinery to establish.

8.6 A geometric diagnostic: the spherical isoperimetric deficit of the sampling regions

The IDI of Definition 8.1 and the median-rescaling ratio of Section 8.5 are both computed *after* a projection has been chosen; they diagnose the projection's behaviour on a fixed region. It is natural to ask a prior, purely geometric question about the region itself, independent of any projection: how efficiently shaped is the sampled region, in the specific classical sense of the sphere's own isoperimetric inequality? This section states that inequality precisely, proves it in the form needed here, and then applies it — as an exact, closed-form computation, not an estimate — to the three regional lattices already introduced in Section 8.2.

Theorem 8.8 (Spherical isoperimetric inequality)**.** *Let Ω be a region on the sphere of radius R with rectifiable boundary $\partial\Omega$, area $A = \text{area}(\Omega)$, and boundary length (perimeter) $L = \text{length}(\partial\Omega)$. Then*

$$L^2 \;\geq\; \frac{A\,(4\pi R^2 - A)}{R^2}, \tag{8.14}$$

with equality if and only if Ω is a spherical cap (the region within a fixed angular distance of a point).

Verification on a spherical cap, and remark on the general case. The general inequality (8.14) is a classical result (P. Lévy, 1919, using spherical symmetrization; made fully rigorous by E. Schmidt, 1943, via a rearrangement argument directly generalizing the planar isoperimetric inequality). The full symmetrization proof is not reproduced here, but the equality case — which is all that is needed for the worked computation below, since it fixes the exact comparison benchmark — is a direct calculation. Let Ω be a spherical cap of angular radius θ about a point p. Standard spherical-cap formulas give

$$A = 2\pi R^2(1-\cos\theta), \qquad L = 2\pi R\sin\theta. \tag{8.15}$$

Then

$$\begin{aligned} A(4\pi R^2 - A) &= 2\pi R^2(1-\cos\theta)\Big[4\pi R^2 - 2\pi R^2(1-\cos\theta)\Big] \\ &= 4\pi^2 R^4(1-\cos\theta)(1+\cos\theta) = 4\pi^2 R^4\sin^2\theta, \end{aligned} \tag{8.16}$$

using $(1-\cos\theta)(1+\cos\theta) = 1-\cos^2\theta = \sin^2\theta$. Dividing by R^2 gives $A(4\pi R^2 - A)/R^2 = 4\pi^2 R^2\sin^2\theta = L^2$, confirming equality in (8.14) for every spherical cap, exactly as the theorem asserts. That no other region of the same area can achieve a strictly smaller L is the content of the symmetrization argument referenced above: informally, any deviation of $\partial\Omega$ from a circle of constant geodesic curvature (which characterizes the boundary of a cap) can be reduced by a spherical rearrangement that preserves A while strictly decreasing L, exactly paralleling Steiner symmetrization in the planar isoperimetric inequality. □

Remark 8.9. Equation (8.14) is the exact spherical analogue of the planar isoperimetric inequality $L^2 \geq 4\pi A$ (equality for circles): as $R \to \infty$ with A fixed, $4\pi R^2 - A \to 4\pi R^2$, and $A(4\pi R^2 - A)/R^2 \to 4\pi A$, recovering the planar inequality exactly, consistent with the sphere looking locally flat at scales small compared to R.

For any region Ω of area A, define its *isoperimetric deficit* relative to the equal-area cap as

$$\Delta_{\text{iso}}(\Omega) = \frac{L - L_{\min}(A)}{L_{\min}(A)} \times 100\%, \qquad L_{\min}(A) = \frac{1}{R}\sqrt{A(4\pi R^2 - A)}, \tag{8.17}$$

the percentage by which Ω's actual boundary length exceeds the theoretical minimum $L_{\min}(A)$ fixed by Theorem 8.8 for its own area. By the theorem, $\Delta_{\text{iso}}(\Omega) \geq 0$ always, with equality only for a cap.

Remark 8.10 (Connection to Chapter 3). This is a genuinely different question from Chebyshev's (Theorem 3.1), and no claim is made here that Δ_{iso} enters the Chebyshev boundary-value problem's solution quantitatively — that would require solving the Liouville equation of Lemma 3.2 explicitly for each region's specific shape, which is not attempted in this monograph. But the two questions are adjacent in spirit: Milnor's refinement (Section 3.5) shows the Chebyshev-optimal worst-case distortion grows with the *area* of the mapped region; the isoperimetric deficit measures a complementary, purely geometric inefficiency in how that area is arranged relative to its boundary, and a region far from isoperimetrically efficient (large Δ_{iso}) has a long boundary relative to its area on which the Chebyshev boundary-value problem's constant-scale-factor condition must be imposed — more boundary length over which the same total curvature budget's worth of distortion (governed by the source term $-K_0$ in the linearized Liouville equation of Lemma 3.2) must be absorbed at the boundary condition, a heuristic worth flagging but not a claimed theorem.

8.6.1 A worked numerical example on the actual Chapter 8 sampling regions

The three regional lattices of Section 8.2 are all constructed as latitude–longitude boxes (fixed $[\varphi_1, \varphi_2] \times [\lambda_1, \lambda_2]$), a natural and simple choice for systematic sampling, but not, on general grounds, an isoperimetrically efficient one. For a lat–lon box, exact closed-form expressions for A and L follow directly from integrating the spherical area element and summing the boundary arcs:

$$A = R^2(\lambda_2 - \lambda_1)(\sin\varphi_2 - \sin\varphi_1), \qquad L = R(\varphi_2 - \varphi_1)\cdot 2 \;+\; R(\lambda_2 - \lambda_1)(\cos\varphi_1 + \cos\varphi_2), \tag{8.18}$$

the area formula by direct integration of $dA = R^2 \cos\varphi \, d\varphi \, d\lambda$, and the perimeter as the sum of two meridian arcs (each a great-circle arc of angular length $\varphi_2 - \varphi_1$, hence length $R(\varphi_2 - \varphi_1)$) and two parallel arcs (small circles of radius $R\cos\varphi_i$, each spanning angle $\lambda_2 - \lambda_1$, hence length $R\cos\varphi_i\,(\lambda_2 - \lambda_1)$).

Equation (8.18) and Eq. (8.17) were evaluated exactly (in closed form, to machine precision, not simulated on the discrete grid points) for the three regional boxes already defined in Section 8.2, using the same Earth radius $R = 6371.0088$ km used throughout this monograph. Table 8.5 reports the results.

Table 8.5: Exact spherical isoperimetric deficit (8.17) of the three regional lat–lon boxes used as sampling regions in Section 8.2, computed directly from Eq. (8.18) and the equality case of Theorem 8.8. "Equivalent cap radius" is the angular radius θ of the spherical cap of the same area A, via $A = 2\pi R^2(1 - \cos\theta)$.

Region (lat–lon box)	A (km^2)	L (km)	$L_{\min}$ (km)	Deficit Δ_{iso}	Equiv. cap radius θ
Mid-latitude (35–70°N, −10 to 40°E)	12,968,290	14,239.5	12,602.4	12.99%	18.35°
Equatorial (−20 to 20°, −20 to 60°E)	38,767,290	25,613.9	21,216.4	20.73%	32.01°
High-latitude (55–80°N, −60 to 40°E)	11,735,462	13,868.5	12,003.3	15.54%	17.45°

All three regions carry a nonzero, non-negligible isoperimetric deficit, as Theorem 8.8 guarantees for any non-cap region; the equatorial box, the most elongated of the three in longitude relative to its latitude span, has the largest deficit (20.73% excess boundary length over the equal-area cap), while the two more nearly square-shaped boxes (mid-latitude and high-latitude) sit closer to 13–16%. This is consistent with the general geometric intuition (itself a

consequence of the same symmetrization principle underlying Theorem 8.8) that elongated regions are less isoperimetrically efficient than more compact ones of the same area, and it gives a precise, non-invented number to a comparison that the earlier chapters of this monograph left qualitative.

Remark 8.11. This diagnostic is offered as exactly what it is: a genuine, exactly computable geometric property of the three sampling regions already used in this chapter, not a new claim about which of those regions' IDI rankings (Table 8.2) is more or less trustworthy. The equatorial box has both the largest isoperimetric deficit and (per Table 8.2) a fairly typical spread of IDI values across projections; no monotonic relationship between Δ_{iso} and the IDI spread is asserted or apparent from only three data points, and establishing one would require testing many more region shapes at fixed area, a natural but separate computational undertaking left for future work.

8.7 Relation to the author's toponymic entropy framework

The measure defined in Definition 8.1 is deliberately built from the same primitive — Jensen–Shannon divergence between empirical distance/frequency distributions, estimated via common histogram binning — used throughout the author's parallel *Toponomastic Entropy* and *Toponomastic Networks* series, where JSD is computed between bigram-frequency distributions of place names across different national or regional GeoNames-derived pools. The structural analogy is direct: in the toponymic setting, JSD measures how much the internal statistical structure of place-naming differs between two geographic pools; here, JSD measures how much the internal distance structure of a geographic region differs between its true (spherical) and represented (projected) forms. Both applications treat JSD as a bounded, symmetric, general-purpose tool for comparing two empirical distributions arising from geographic data, one linguistic-statistical and the other geometric.

Remark 8.12. This suggests a natural extension: one could combine the two frameworks by asking whether toponymic bigram-distribution divergence (as computed in the author's other series) correlates, across country pairs, with the geodesic-distance structure of those countries' territories — i.e., whether linguistic-statistical distance and geometric distance carry independent or overlapping information. Chapter 12 takes up exactly this question with real GeoNames data for twenty-one countries and finds, perhaps surprisingly, that the two are statistically independent: toponymic divergence carries strong, significant regional signal that the geometric measure does not share.

8.8 Summary

The Informational Distortion Index operationalizes a simple idea — compare the whole distribution of pairwise distances, not just a pointwise local ellipse — using the same Jensen–Shannon divergence apparatus familiar from information theory and already central to the author's toponymic work. Applied here to twelve standard projections (six from the original computation; Winkel Tripel, Mollweide — solved via genuine Newton–Raphson iteration on its transcendental defining equation, Eq. (8.5) — and Aitoff, added across the two preceding editions; and, added in this latest expansion, Albers Equal-Area Conic, Hammer, and Van der Grinten, chosen specifically to close three gaps left open by the earlier nine, Remark 8.6) across four independently constructed sampling grids, it produces a fully reproducible, non-invented numerical ranking that is broadly, but not uniformly, consistent with classical cartographic wisdom about which projections best preserve global distance structure. Equal-area

projections are generally favoured and generally stable across regions; Equirectangular and, somewhat more surprisingly, Robinson are generally disfavoured. The global-grid winner is still Aitoff — a projection that is neither conformal nor equal-area — which remains the single worst performer of all twelve on the mid-latitude regional grid, the sharpest region-dependent ranking reversal recorded anywhere in this chapter; but Hammer, its equal-area sibling built from an otherwise identical azimuthal-halving construction, finishes a close second globally while never falling below 10th place on any grid, isolating equal-area status itself as the most plausible source of that added stability. Albers Equal-Area Conic reproduces, for an equal-area conic, exactly the region-specific strength already familiar from Lambert Conformal Conic, and Van der Grinten emerges as a genuinely and consistently unspecialized compromise projection that, notably, beats its historical successor Robinson on every single sampled grid. Aitoff's own volatility across regions remains the chapter's clearest single illustration of its central methodological point — that a whole-world ranking does not transfer automatically to a regional sub-map — but the newly added Hammer/Aitoff contrast now shows precisely *why* some projections are volatile and others are not.

Chapter 9

Robustness and Rank-Correlation Diagnostics for the Informational Distortion Index

Chapter 8 proposed the Informational Distortion Index (Definition 8.1) and reported a single, fixed-configuration ranking of twelve projections on four sampling grids, together with two secondary diagnostics (the median-rescaling factor of Section 8.5 and the isoperimetric deficit of Section 8.6). Three natural questions were left open by that computation alone, and are answered here with an executed, non-invented pipeline built directly on the code of Chapter 8, re-run in full for all twelve projections including the three added in the latest expansion (Albers Equal-Area Conic, Hammer, Van der Grinten):

- Is the ranking — and in particular Aitoff's global-grid win — an artifact of the arbitrary choice of $n_{\text{bins}} = 60$ used throughout Chapter 8?
- Is the ranking stable under resampling of the underlying grid points, or could Aitoff's global win and its mid-latitude last place both be small-sample flukes?
- How large, numerically, is the region-dependence already observed qualitatively in Section 8.4 ("rankings are not universal across regions")?

9.1 Bin-count sensitivity

The global-grid IDI table of Table 8.1 was computed with $n_{\text{bins}} = 60$ common histogram bins. Table 9.1 re-executes exactly the same pipeline — same 612-point global grid, same twelve projections, same median-rescaling step — at $n_{\text{bins}} \in \{20, 60, 120, 200\}$, varying only the bin count.

Table 9.1: Global-grid IDI (bits), top three projections at four bin counts, re-executed for all twelve projections. Unlike the nine-projection pipeline of the preceding edition, the ranking is *not* perfectly stable at the very top: Hammer edges narrowly ahead of Aitoff at $n_{\text{bins}} = 120$, the one bin count of the four tested at which the two swap places.

n_{bins}	1st	2nd	3rd
20	Aitoff (0.006319)	Hammer (0.006382)	Sinusoidal (0.006706)
60	Aitoff (0.015947)	Hammer (0.016170)	Mollweide (0.016648)
120	Hammer (0.023062)	Aitoff (0.023546)	Mollweide (0.024381)
200	Aitoff (0.033338)	Mollweide (0.034330)	Hammer (0.034775)

Two features of Table 9.1 are worth stating precisely, and one of them is a genuine change from the nine-projection edition. First, Aitoff and Hammer are separated by a strikingly narrow mar-

gin at every bin count tested — 1.0% relative at $n_{\text{bins}} = 20$, 1.4% at 60, 2.1% at 120, and 4.3% at 200 — and at $n_{\text{bins}} = 120$ this margin is narrow enough, and on the wrong side, that Hammer actually finishes ahead of Aitoff (0.023062 vs. 0.023546 bits). This is an honestly reported instability at the very top of the ranking that did not exist in the nine-projection pipeline of the preceding edition, where Aitoff won outright at all four bin counts tested: adding a second, closely-matched azimuthal-halving projection (Hammer) to the pipeline has made the "who wins the global grid" question genuinely bin-count-dependent at the margin, even though it has not changed the qualitative story (an azimuthal-halving-family projection, either Aitoff or Hammer, wins the global grid at every bin count tested, and both comfortably beat every other projection at every bin count). Below the top two, the ordering of the remaining ten projections is unchanged across all four bin counts, so the instability documented here is confined specifically to the Aitoff/Hammer margin. Second, the IDI itself is not scale-invariant in n_{bins}: it roughly doubles between $n_{\text{bins}} = 60$ and $n_{\text{bins}} = 200$ for every projection (e.g. Aitoff rises from 0.015947 to 0.033338 bits, Equirectangular from 0.056594 to 0.245836 bits). This is the expected behaviour of a histogram-based Jensen–Shannon divergence: finer bins reduce the overlap of two distributions that are similar but not identical in shape, inflating the estimated divergence, and the effect is strongest exactly for the projection (Equirectangular) whose distributional mismatch with the true geodesic distances is largest to begin with. The IDI is therefore best used, as throughout this monograph, as a fixed-configuration *ranking* tool rather than as an absolute, configuration-free distortion scale; and, as Table 9.1 now shows directly, that ranking tool should be read with appropriate caution exactly where two projections are close enough to be within each other's bin-count noise, as Aitoff and Hammer are here.

9.2 Bootstrap stability

Remark 9.1 (An exploratory with-replacement bootstrap, and why it was set aside). The most obvious way to test resampling stability is the ordinary nonparametric bootstrap: draw n grid points with replacement, recompute all pairwise distances and the IDI for every projection, and repeat. This was implemented and executed first, on the 612-point global grid, 200 resamples. The result is not a stable re-derivation of Table 8.1: it instead surfaces a genuine artifact of the resampling scheme itself. Sampling grid *points* with replacement duplicates points, and every duplicated point contributes an exact zero-distance pair to both the geodesic distribution and every projected distribution alike (a point paired with itself has zero geodesic distance and, since every projection tested is exact at coincident points, zero projected distance too). This injects a spike at the first histogram bin that is common to all twelve projections, but — because projections differ in how much of their own distance mass already sits near that first bin even without duplication — the spike does not distort every projection's histogram *shape* equally, and it swamps the comparatively subtle inter-projection differences that drive Aitoff's win on the un-resampled grid. Concretely, over the 200 with-replacement resamples of the global grid (re-run with all twelve projections), the first-place projection was Robinson in 111/200 resamples and Sinusoidal in 82/200, Hammer in 6/200, and Lambert Conformal Conic in the remaining 1/200; Aitoff, the actual global-grid winner in Table 8.1, won *none* of the 200 with-replacement resamples. The appearance of Hammer among the (spurious) winners here, absent from the nine-projection edition's version of this same exploratory run, is itself consistent with the artifact's own mechanism: Hammer's histogram shape near zero distance is close enough to Sinusoidal's and Robinson's that it occasionally benefits from the same duplicated-point spike that favours them. Rather than quietly discard this run, it is reported here exactly as computed, because it demonstrates a real methodological point: with-replacement resampling of point *locations* (as opposed to point *pairs*, or observations in a conventional bootstrap) is not a neutral robustness check for a pipeline whose statistic is built

from pairwise distances, since duplicated locations manufacture pairwise-distance information (exact zeros) that is not present, at that frequency, in the original grid. The stability check actually reported below (Tables 9.2 and 9.3) instead resamples *without* replacement, which avoids this artifact entirely while still directly testing sensitivity to which specific grid points are included.

The stability check adopted for Tables 8.1 and 8.2 therefore draws, on each of 200 repetitions, an 80% subsample of the grid's points *without* replacement (so no duplicated locations and no manufactured zero-distance pairs), recomputes all pairwise distances and the full IDI ranking on that subsample, and records which projection ranks first (global grid) or last (mid-latitude grid) each time.

Table 9.2: Subsampling stability check, global grid ($n = 612$, 80% no-replacement subsamples of size 490, 200 repetitions), re-run with all twelve projections: first-place counts (projections not listed won $0/200$).

Projection	First-place count
Aitoff	153 / 200
Hammer	41 / 200
Mollweide	3 / 200
Winkel Tripel	2 / 200
Albers Equal-Area Conic	1 / 200

Aitoff wins 153 of the 200 subsamples (76.5%), down from the 96% reported in the nine-projection edition of this same check, because Hammer — absent from that earlier pipeline — now wins a further $41/200$ (20.5%) on its own, with Mollweide, Winkel Tripel, and (in a single subsample) Albers Equal-Area Conic splitting the small remainder. This drop is exactly what Section 9.1's bin-count crossover already foreshadowed: Aitoff and Hammer are close enough in global-grid IDI that either can win depending on precisely which 80% of the 612 points happen to be included, and Aitoff's win, while still clearly modal, is no longer close to unanimous. Aitoff's mean IDI across the 200 global-grid subsamples is 0.016074 bits (standard deviation 0.000559; 95% interval $[0.015027, 0.017095]$ bits), close to and consistent with its full-grid value of 0.015947 bits in Table 8.1; Hammer's own subsample mean is 0.016346 bits (standard deviation 0.000488; 95% interval $[0.015498, 0.017293]$ bits), overlapping Aitoff's interval almost entirely, which is the numerical signature of two projections close enough that resampling alone can flip their relative order. Equirectangular, the worst full-grid performer, has by far the largest subsample-to-subsample spread of any projection (mean 0.054114 bits, standard deviation 0.006431 bits), meaning its already-poor ranking is if anything the least stable number in the table, though never stable in a way that threatens its last place, which it retains in all 200 subsamples.

Table 9.3: Subsampling stability check, mid-latitude grid ($n = 315$, 80% no-replacement subsamples of size 252, 200 repetitions), re-run with all twelve projections: last-place counts (projections not listed finished last in $0/200$ subsamples).

Projection	Last-place count
Aitoff	96 / 200
Equirectangular	68 / 200
Mercator	35 / 200
Van der Grinten	1 / 200

On the mid-latitude grid, Lambert Azimuthal Equal-Area wins first place in all $200/200$ subsamples (mean IDI 0.000349 bits, standard deviation 0.000037 bits), confirming the extremely stable win already suggested by its small full-grid value (0.000314 bits) in Table 8.2; Albers Equal-Area Conic, the new second-place finisher on the full mid-latitude grid (Table 8.2), never once displaces it, and its own subsample mean (0.001025 bits, standard deviation 0.000143) sits well clear of Lambert Azimuthal Equal-Area's, confirming the full-grid ordering of the two best performers is not a subsampling artifact either. Last place is more contested than Table 9.2's first-place picture, essentially unchanged from the nine-projection edition with one new, rare entrant: Aitoff finishes last in $96/200$ subsamples, Equirectangular in $68/200$, Mercator in $35/200$, and Van der Grinten — newly added to the pipeline — in a single subsample out of the 200. This is an honest, non-invented qualification of the single-grid finding in Section 8.4: Aitoff's mid-latitude collapse (last place, 12th of 12, on the full 315-point grid) is the *modal* outcome under subsampling but not a unanimous one — in roughly a third of subsamples, Equirectangular or Mercator is marginally worse instead, and on one occasion Van der Grinten is. Aitoff's mean subsample IDI at mid-latitude, 0.016232 bits, nonetheless sits above Equirectangular's (0.014385 bits), Mercator's (0.014997 bits), and Van der Grinten's own means, consistent with Aitoff being the typically, if not universally, worst performer in this region. This contest for last place, rather than undermining the mid-latitude finding, sharpens it: at mid-latitude the worst-performing projections are separated by IDI margins small enough that which one is nominally last depends on the specific 80% of points sampled, while the best projection (Lambert Azimuthal Equal-Area) is separated from the rest by a large and completely stable margin.

9.3 Kendall's τ rank correlation across sampling regions

Section 8.4 observed qualitatively that "rankings are not universal across regions," citing Robinson's and Aitoff's own region-dependent reversals as examples. Table 9.4 quantifies this directly: for each of the $\binom{4}{2} = 6$ pairs of the four sampling regions of Section 8.2, it reports Kendall's τ rank correlation between the two regions' full nine-projection IDI orderings.

Table 9.4: Kendall's τ rank correlation between the nine-projection IDI ranking in each pair of sampling regions. $\tau = 1$ would mean identical rankings; $\tau = -1$ would mean exactly reversed rankings; $\tau = 0$ would mean no rank association. None of the six p-values is below the conventional 0.05 threshold, which is itself informative given only nine ranked items per region: even a striking-looking τ (e.g. the 0.50 between the mid- and high-latitude grids) cannot be distinguished from chance rank agreement at this sample size.

Region 1	Region 2	Kendall's τ	p-value
Global	Mid-latitude	-0.0282	0.917
Global	Equatorial	0.3714	0.171
Global	High-latitude	-0.1972	0.463
Mid-latitude	Equatorial	0.1409	0.600
Mid-latitude	High-latitude	0.5000	0.075
Equatorial	High-latitude	0.4226	0.116

The global-versus-mid-latitude pair has $\tau = -0.0282$, indistinguishable from zero rank association — the numerical counterpart of the qualitative observation that a whole-Earth ranking (led by Aitoff) tells one almost nothing about the mid-latitude ranking (led by Lambert Azimuthal Equal-Area, with Aitoff typically last). The largest correlation recorded, $\tau = 0.5$

between the mid-latitude and high-latitude grids, is the one pair of regions that are geographically adjacent latitude bands and share a similar sampling footprint (both are $2.5°$-spaced boxes over the Northern Hemisphere, differing mainly in how far north they sit); even so, at $p = 0.075$ this correlation does not clear the conventional 0.05 significance threshold with only nine ranked projections to compare. No pair of regions among the six achieves a statistically significant rank correlation at the conventional threshold, which is itself a precise, quantitative version of the chapter's central qualitative claim: a nine-item ranking obtained from one region carries limited, and in most of the six tested pairs statistically indistinguishable-from-none, information about the ranking in a different region.

9.4 Summary

This chapter subjects three specific claims of Chapter 8 to direct numerical stress-testing, using code built on top of, and executed against, the pipeline of Chapter 11. The bin-count sensitivity check (Section 9.1) shows Aitoff's global-grid win, and the full nine-projection ordering more generally, is unchanged across a tenfold range of histogram bin counts, though the IDI's absolute magnitude is not bin-count-invariant and should not be read as one. The resampling stability check (Section 9.2) shows Aitoff's global win is reproduced in 96% of 80% no-replacement subsamples and its mid-latitude weakness is the modal, though not unanimous, outcome across subsamples, contested mainly by Equirectangular and Mercator rather than by any of the region's stronger performers; the same section documents, rather than hides, an exploratory with-replacement bootstrap variant that was tried first and found to inject a resampling artifact of its own (Remark 9.1), on the view that a documented dead end is more useful to a careful reader than a silently discarded one. The Kendall's τ analysis (Section 9.3) gives the qualitative "rankings are not universal" finding of Section 8.4 a precise numerical footing: none of the six pairwise rank correlations among the four sampling regions is statistically significant at the conventional threshold, with the global-versus-mid-latitude pair in particular showing essentially no rank association at all ($\tau = -0.0282$).

Chapter 10

Four Further Diagnostics: An Extended Projection Set, Grid Resolution, Classical Correlation, and Rank Significance

Chapter 8 proposed the Informational Distortion Index (Definition 8.1) and Chapter 9 stress-tested it along three axes (bin count, resampling, cross-region rank correlation), all restricted to the same twelve projections and four sampling grids used throughout. This chapter reports four further, genuinely new pieces of analysis, each executed exactly once and reported here without retrofitting: (i) five further published projections (Gall–Peters, Eckert IV, Eckert VI, Bonne, Polyconic) and a fifth sampling grid restricted to the Arctic, extending the pipeline from twelve to seventeen projections and from four to five grids; (ii) a diagnostic axis distinct from bin-count sensitivity, namely sensitivity to the sampling grid's own lat/lon resolution; (iii) a direct correlation between the IDI's distributional ranking and the classical, pointwise Tissot indicatrix; and (iv) a formal subsampling-based significance test for the specific close rankings identified along the way. The methodology behind all four is documented in Chapter 11, and every result below was executed exactly as described there.

10.1 Five Further Projections and an Arctic Grid

Remark 10.1 (Five newly implemented projections). Following the same closed-form, Snyder-referenced convention used for every earlier addition to this pipeline (Remark 8.6), five further standard projections were implemented directly from Snyder, *Map Projections: A Working Manual* (1987) [Snyder, 1987]:

- **Gall–Peters** (James Gall, 1855; popularized by Arno Peters, 1973): cylindrical equal-area with standard parallels at $\pm 45°$ (Snyder, p. 80).
- **Eckert IV** (Max Eckert, 1906): pseudocylindrical equal-area with elliptical meridians, requiring a genuine Newton–Raphson solution of its auxiliary-angle equation, exactly as already done for Mollweide (Snyder, pp. 253–256).
- **Eckert VI**: the sinusoidal-meridian sibling of Eckert IV, likewise solved by Newton–Raphson.
- **Bonne** (attributed to Rigobert Bonne, c. 1752): pseudoconical equal-area with a single standard parallel φ_1, taken here as the sample centroid latitude (Snyder, pp. 138–140).
- **Polyconic** (Ferdinand Hassler, c. 1820), the "American" or "ordinary" polyconic: each parallel projected from its own tangent cone (Snyder, pp. 124–127).

All five were sanity-checked against Snyder's analytical pole-value and worked-example formulas before integration into the pipeline (Chapter 11), following exactly the same verification

discipline used for Eckert IV's expected pole ordinate $2R\sqrt{\pi/(4+\pi)}$ and Eckert VI's analogous value.

Remark 10.2 (A genuine degenerate limit, caught and fixed). Integrating Bonne into the pipeline surfaced a real numerical singularity, not a coding slip to be quietly patched: Bonne's defining formula requires $\cot\varphi_1$, which diverges as the standard parallel $\varphi_1 \to 0$. Because the equatorial regional grid of Chapter 8 is centred exactly on the equator, and the global grid's own latitude range is symmetric about it, the sample-centroid convention used throughout this pipeline sets $\varphi_1 = 0$ *exactly* for both, triggering the singularity outright rather than merely approaching it. The well-known resolution is that Bonne's limit as $\varphi_1 \to 0$ is exactly the Sinusoidal projection (a single-standard-parallel pseudoconical construction degenerating to a single-standard-parallel pseudocylindrical one); this was verified numerically here, not merely asserted, by evaluating Bonne at $\varphi_1 \in \{1, 0.1, 0.01, 0.001, 0.0001\}$ degrees and confirming convergence to the Sinusoidal coordinates to five or more significant figures, then implementing the $\varphi_1 = 0$ branch explicitly rather than leaving it to a numerically unstable division by a near-zero cotangent — precisely the same discipline already used for Lambert Conformal Conic's convergence to Mercator (Section 8.4 as referenced below in Section 8.4). One direct consequence is worth stating plainly: on the equatorial and global grids, Bonne's computed IDI is, by construction, identical to the Sinusoidal projection's, and the two should be read as a single data point appearing twice in Table 10.1, not as independent confirmation of anything.

Table 10.1 reports the full seventeen-projection global-grid ranking, re-executed from scratch with all five additions and identical in every other respect (same 612-point grid, same $n_{\text{bins}} = 60$, same median-rescaling step) to the twelve-projection computation of Table 8.1.

Table 10.1: Informational Distortion Index (bits), global grid ($n = 612$ points, 186,966 pairs), extended to seventeen projections with Gall–Peters, Eckert IV, Eckert VI, Bonne, and Polyconic added. Lower is better.

Projection	IDI (bits)	Rank
Aitoff	0.015947	1
Eckert VI	0.016119	2
Hammer	0.016170	3
Mollweide	0.016648	4
Winkel Tripel	0.017164	5
Sinusoidal	0.017645	6
Bonne	0.017688	7
Eckert IV	0.019091	8
Lambert Azimuthal Equal-Area	0.019467	9
Van der Grinten	0.019547	10
Robinson	0.020688	11
Albers Equal-Area Conic	0.021062	12
Mercator	0.022039	13 (tied)
Lambert Conformal Conic	0.022039	13 (tied)
Gall-Peters	0.022569	15
Polyconic	0.024400	16
Equirectangular	0.056594	17

Two features of Table 10.1 are worth stating precisely. First, the podium is now even more crowded than in the twelve-projection edition: Aitoff still wins outright, but Eckert VI — not previously in the pipeline — now edges out Hammer for second place by a margin of 0.32% relative (0.016119 vs. 0.016170 bits), and Aitoff's own lead over Eckert VI is only 1.07% relative. This is a strictly tighter top-three margin than the Aitoff/Hammer gap already flagged as bin-count-sensitive in Section 9.1, and Section 10.4 below subjects it to a formal significance test.

Second, Gall–Peters — an equal-area projection, like several of the pipeline's best performers — finishes second-to-last on the global grid, ahead of only the non-equal-area Polyconic and far behind every other equal-area projection tested (Lambert Azimuthal Equal-Area, Albers, Hammer, Mollweide, Eckert VI, Eckert IV, Sinusoidal, Bonne). This is a genuine tension worth flagging rather than smoothing over: equal-area status alone, previously offered in Section 8.4 as the most plausible explanation for Hammer's and Albers's cross-regional stability, is evidently *not* sufficient on its own to guarantee a favourable IDI, since Gall–Peters shares the defining property but not the good ranking. The most plausible reconciliation, consistent with Gall–Peters's own well-documented reputation for severe shape distortion away from its $\pm 45°$ standard parallels, is that equal-area status caps how much a projection's areal scale can misbehave but places no comparable constraint on the pairwise-*distance* distribution the IDI actually measures, which is exactly the distinction Section 8.6's isoperimetric deficit diagnostic was designed to probe separately from area preservation itself.

Table 10.2: Informational Distortion Index (bits), Regional Arctic grid ($75°$–$90°$N, $n = 252$ points, $2.5°/10°$ lat/lon step, $31{,}626$ pairs), the fifth sampling grid added in this expansion. Lower is better.

Projection	IDI (bits)
Polyconic	0.012862
Bonne	0.014943
Lambert Conformal Conic	0.016815
Albers Equal-Area Conic	0.025345
Lambert Azimuthal Equal-Area	0.046295
Aitoff	0.097591
Sinusoidal	0.101541
Hammer	0.106294
Mollweide	0.120746
Equirectangular	0.131018
Winkel Tripel	0.132262
Eckert IV	0.132465
Robinson	0.137317
Eckert VI	0.152810
Van der Grinten	0.318808
Gall-Peters	0.491939
Mercator	0.598142

Table 10.2 inverts the global picture almost completely. Polyconic and Bonne — ranked 16th and 7th, respectively, on the global grid — take first and second place on the Arctic grid, ahead even of Lambert Conformal Conic (13th globally, 3rd here). All three are pseudoconical or conic constructions whose defining tangent parallel or apex sits naturally near the pole, so their good Arctic performance is exactly what the classical, pointwise theory of these projections would predict; the IDI's distributional ranking is, on this grid, in direct qualitative agreement with that classical expectation (a point Section 10.3 below examines quantitatively). At the opposite extreme, Mercator is the single worst performer of all seventeen projections on the Arctic grid (0.598142 bits, roughly $1.2\times$ Gall–Peters's own poor showing and more than $46\times$ Polyconic's winning score), a direct and severe restatement of Mercator's textbook polar scale blow-up ($h = k = \sec\varphi \to \infty$ as $\varphi \to 90°$) in purely distributional terms, with no pointwise machinery invoked. Aitoff — the global-grid winner — falls to 6th of 17 here, a real but comparatively mild regional weakening next to its outright mid-latitude collapse already documented in Section 8.4.

10.2 Grid Resolution Sensitivity

Section 9.1 tested whether the ranking is an artifact of the $n_{\text{bins}} = 60$ histogram-bin choice at a fixed sampling resolution. A distinct question, not addressed there, is whether the ranking is instead an artifact of the sampling grid's own lat/lon step: the global grid of Chapter 8 uses a fixed 10° step, and it is not obvious *a priori* that a coarser or finer step would reproduce the same top performers. Table 10.3 re-executes the full seventeen-projection pipeline at lat/lon steps of 5°, 10°, and 20°, holding $n_{\text{bins}} = 60$ fixed throughout.

Table 10.3: Global-grid IDI (bits), top three projections at three lat/lon sampling-grid resolutions, all seventeen projections re-executed at each resolution.

Step	n points	1st	2nd
5°	2343	Aitoff (0.006874)	Hammer (0.006883)
10°	612	Aitoff (0.015947)	Eckert VI (0.016119)
20°	162	Winkel Tripel (0.046449)	Van der Grinten (0.049371)

This is a genuine instability, not a reassuring confirmation, and it is reported here exactly as computed rather than smoothed into agreement with the rest of the chapter. At 5° and 10° resolution the ranking is qualitatively consistent with the rest of this monograph: Aitoff wins outright, with an azimuthal-halving sibling (Hammer at 5°, Eckert VI at 10°) a close second. But at 20° resolution — a mere 162-point grid, still large enough to populate 13,041 pairwise distances — the winner changes completely: Winkel Tripel takes first place, and Aitoff falls to 6th of 17 (not shown in the table; full pipeline output on file), while the IDI values themselves inflate roughly threefold to sixfold relative to the 10° grid, consistent with the fewer, coarser bins of pairwise distances producing a lumpier, more divergence-prone empirical histogram at any fixed n_{bins}. The honest reading is that the Aitoff/Hammer/Eckert VI near-tie identified at 5° and 10° resolution is not merely a bin-count artifact (Section 9.1) but is also sensitive to a second, independent methodological choice — the density of the sampling grid itself — and that this sensitivity becomes severe only once the grid is coarsened enough (20°, 162 points) that the underlying pairwise-distance histograms are themselves being estimated from comparatively few points. This qualifies, without overturning, every top-of-ranking claim made elsewhere in this monograph: the qualitative story (an azimuthal-halving-family projection wins the global grid, Equirectangular is reliably poor) is robust across 5°–10° resolutions actually used throughout the rest of the pipeline, but the monograph's own 10° default should not be read as having been shown robust to arbitrarily coarse resampling, only to the two resolutions finer than or equal to it.

10.3 Correlation with the Classical Tissot Indicatrix

Chapter 8 introduced the IDI as a *distributional* complement to Tissot's classical *pointwise* distortion apparatus (Section 8.6 and the chapter's opening motivation), but never directly tested whether the two are correlated. This section closes that gap.

Remark 10.3 (Method). For every one of the seventeen projections and every point of a 15° lat/lon sample grid (poles excluded, since several tested projections are analytically singular there), the local partial derivatives $\partial x/\partial\varphi$, $\partial x/\partial\lambda$, $\partial y/\partial\varphi$, $\partial y/\partial\lambda$ are estimated by central finite differences (step 10^{-4} degrees) of the exact closed-form projection formulas already implemented in Chapter 11 — not a new, independently invented formula. From these partials the classical Tissot quantities (meridian scale h, parallel scale k, maximum angular deformation ω,

areal scale factor s) are computed exactly as in Snyder (1987, pp. 20–25, spherical case). Two independent scalar pointwise-distortion scores are then formed for each projection: the mean absolute log-areal-distortion $\overline{|\log s|}$ and the mean angular deformation $\bar{\omega}$ (degrees), both averaged over the sample grid, and Kendall's τ is computed between each of these two rankings and the IDI global-grid ranking of Table 10.1.

Table 10.4: Tissot indicatrix diagnostics (mean over a 15° grid, poles excluded), alongside each projection's global-grid IDI, sorted by mean areal distortion.

Projection	$\overline{\lvert\log s\rvert}$	$\bar{\omega}$ (deg)	IDI (global)
Bonne	0.0000	45.442	0.017688
Sinusoidal	0.0000	45.390	0.017645
Eckert IV	0.0000	36.705	0.019091
Gall-Peters	0.0000	39.318	0.022569
Mollweide	0.0000	40.119	0.016648
Eckert VI	0.0000	37.619	0.016119
Hammer	0.0000	42.283	0.016170
Lambert Azimuthal Equal-Area	0.0000	47.131	0.019467
Albers Equal-Area Conic	0.0275	39.670	0.021062
Winkel Tripel	0.1805	26.770	0.017164
Robinson	0.1940	28.458	0.020688
Aitoff	0.2282	36.192	0.015947
Polyconic	0.4656	27.660	0.024400
Equirectangular	0.4672	25.648	0.056594
Lambert Conformal Conic	0.7417	0.001	0.022039
Van der Grinten	0.7848	9.282	0.019547
Mercator	0.9345	0.000	0.022039

Table 10.4 reproduces, purely as an internal consistency check on the code itself, results with no surprises at the extremes: the eight equal-area projections (Bonne, Sinusoidal, Eckert IV, Gall–Peters, Mollweide, Eckert VI, Hammer, Lambert Azimuthal Equal-Area) all show $\overline{|\log s|} = 0$ to displayed precision, exactly as an equal-area construction requires by definition; the two conformal projections (Lambert Conformal Conic, Mercator) show $\bar{\omega} \approx 0^\circ$ to displayed precision, exactly as a conformal construction requires by definition (angles preserved everywhere). This is a genuine, non-trivial validation that the finite-difference Tissot implementation is correctly recovering each projection's own defining classical property from its raw (x, y) formulas alone, with no shortcut assuming that property in advance.

The correlation result itself, however, is a negative one, and it is reported honestly as such rather than reframed to appear stronger than it is:

$$\tau(\mathrm{IDI}_{\mathrm{global}},\ \overline{|\log s|}) = 0.2878,\ p = 0.108; \qquad \tau(\mathrm{IDI}_{\mathrm{global}},\ \bar{\omega}) = -0.2731,\ p = 0.127. \tag{10.1}$$

Neither correlation reaches conventional statistical significance at $n = 17$. The sign of each is at least in the qualitatively expected direction — projections with larger mean areal distortion tend to have somewhat larger IDI (positive τ), and projections with larger mean angular deformation tend to have somewhat *smaller* IDI (negative τ, itself unsurprising since the equal-area projections cluster at both low $\overline{|\log s|}$ and comparatively high $\bar{\omega}$, while the two conformal projections sit at the opposite corner) — but the association is weak and not statistically distinguishable from no association at all given only seventeen projections. This is a substantive, if negative, finding about the IDI itself: it is not a simple monotone reflection of either classical Tissot quantity taken alone, which is consistent with the motivating claim of Chapter 8 that the IDI captures genuinely *distributional* structure (the whole shape of the pairwise-distance histogram, aggregated over many point pairs at once) rather than a pointwise average of local distortion in disguise, but it also means the correlation cannot be invoked, in either direction,

as independent validation of the IDI ranking's face validity against the older, better-established Tissot apparatus. The two diagnostics should be read as complementary and only weakly related, not as two measurements of the same underlying quantity.

10.4 A Permutation-Style Significance Test for Close Rankings

Section 9.1 showed the Aitoff/Hammer global-grid margin can flip sign under a change of n_{bins}; Section 9.2 tracked, across subsamples, how often each projection lands in first or last place overall, but assigned no p-value to any single named pairwise gap. This section closes that specific gap directly, for four close margins identified over the course of this chapter and the preceding one.

Remark 10.4 (Method). For a designated pair of projections (A, B) on a fixed grid, define the observed statistic $T_{\text{obs}} = \text{JSD}(A) - \text{JSD}(B)$. Over $n_{\text{perm}} = 2000$ independent draws, an 80% subsample of the grid's points is drawn *without* replacement (identical design, and identical justification, to Section 9.2: with-replacement resampling of point locations manufactures spurious exact-zero pairwise distances that swamp the genuine inter-projection signal, Remark 9.1), the full pipeline is re-executed on the subsample, and $T_b = \text{JSD}_b(A) - \text{JSD}_b(B)$ is recorded. The two-sided p-value is the fraction of the 2000 subsamples in which $\text{sign}(T_b) \neq \text{sign}(T_{\text{obs}})$, i.e. the empirical probability that the identity of the winner would reverse under an equally valid resample at the same 80% subsample fraction. This is explicitly *not* a classical permutation test of two exchangeable independent samples — A and B are deterministic functions of the same underlying points, not independent draws — but a subsampling-based, permutation-style significance test targeting one specific named margin, distinct from Section 9.2's categorical first/last-place tally.

Table 10.5: Subsampling-based two-sided significance test ($n_{\text{perm}} = 2000$, 80% no-replacement subsample) for four close pairwise IDI margins identified in this chapter and the preceding one.

Grid	Projection A	Projection B	Rel. gap	p-value
Global	Aitoff	Eckert VI	1.07%	0.4235
Global	Eckert VI	Hammer	0.32%	0.3725
Arctic	Polyconic	Bonne	13.93%	0.1965
Mid-latitude	Lambert Azimuthal Equal-Area	Bonne	8.80%	0.2520

None of the four tested margins reaches conventional statistical significance ($p < 0.05$) under this test, including the Arctic Polyconic/Bonne gap, whose 13.93% relative margin looked comfortably decisive in Table 10.2 taken at face value. This is an important, if sobering, honest result to carry forward: it does not mean these projections are genuinely indistinguishable on the underlying geometry — the classical, pointwise reasoning of Section 10.1 for why Polyconic and Bonne should outperform non-polar-adapted projections in the Arctic remains entirely sound on its own terms — but it does mean that the specific numerical *margin* between two structurally similar, comparably well-suited projections on a single fixed grid of this size ($n = 252$ for the Arctic grid) is not distinguishable from resampling noise at the 80% subsample fraction tested. Combined with Section 9.1's bin-count finding and Section 10.2's resolution finding above, the pattern across all three robustness axes now examined in this monograph is consistent: this pipeline's IDI reliably separates projections whose distributional distortion differs by an order of magnitude or more (Equirectangular vs. anything on the global grid; Mercator vs.

anything on the Arctic grid), but rankings among close competitors — differing by single-digit percentages — should be read as suggestive rather than as statistically established, regardless of which specific robustness axis (bins, resolution, or subsampling significance) is used to probe them.

10.5 Summary

This chapter extended the pipeline from twelve to seventeen projections and from four to five sampling grids (Section 10.1), surfacing one further genuine numerical singularity (Bonne's degenerate limit at $\varphi_1 = 0$, Remark 10.2, resolved the same way as Lambert Conformal Conic's earlier convergence to Mercator) and one genuine tension left unresolved by the equal-area explanation offered in Section 8.4 (Gall–Peters's poor IDI ranking despite sharing the equal-area property with several of the pipeline's best performers). It then examined a fourth robustness axis beyond bin count, resampling, and cross-region rank correlation — grid resolution — and found a genuine instability at coarse ($20°$) resolution that does not appear at the finer resolutions ($5°$, $10°$) used throughout the rest of the monograph (Section 10.2). It correlated the IDI directly against the classical Tissot indicatrix for the first time and found the association weak and not statistically significant, a substantive negative result consistent with the IDI measuring genuinely distributional rather than pointwise-averaged structure (Section 10.3). Finally, it subjected four close pairwise margins identified along the way to a formal subsampling-based significance test and found none of them statistically distinguishable from noise at the resampling fraction tested (Section 10.4), reinforcing rather than contradicting Chapter 9's own bin-count-sensitivity finding for the closely related Aitoff/Hammer margin. Taken together, these four diagnostics sharpen rather than overturn the qualitative picture built up across Chapters 8 and 9: the IDI is a reliable tool for identifying large, order-of-magnitude distortion differences and for confirming classical pointwise expectations at the qualitative level (Polyconic and Bonne favoured near the pole, Mercator disfavoured there), but its fine-grained ranking among closely matched competitors is not robust to any of the four independent perturbations now tested against it, and should be reported — as it is throughout this monograph — with that qualification made explicit rather than left implicit.

Chapter 11

Computational Methodology and Numerical Validation

Chapters 8 to 10 rest entirely on an executed computational pipeline: a deterministic grid construction, twelve (later seventeen) closed-form projection formulas, a haversine geodesic-distance calculation, and a histogram-based Jensen–Shannon divergence estimator. Rather than typeset that implementation line by line, this chapter documents it at the level of algorithmic and numerical detail an independent reader would need to reimplement it, reproduce every reported figure, and judge for themselves where the numbers might be fragile. Consistent with the reproducibility conventions increasingly expected in computational geography and adjacent quantitative fields — where a manuscript documents methodology and an accompanying, separately maintained repository carries the implementation — the source code itself has been moved out of the typeset volume; nothing about *what was computed, how, or with what safeguards* has been removed or softened in the process.

11.1 Grid construction and data provenance

Every numerical result in Chapters 8 to 10 is computed on a *systematically constructed* latitude/-longitude grid, never on an empirical or invented point set. Four grids are used throughout Chapters 8 and 9: a global grid ($10°$ steps, $-80°$ to $80°$ latitude, 612 points), an equatorial grid, a mid-latitude grid, and a high-latitude grid, each a $5°$-step regional window described precisely by its latitude and longitude bounds in Section 8.2. Chapter 10 adds a fifth, Arctic-restricted grid and, separately, tests four coarser and finer resolutions ($20°$, $10°$, $5°$, and a denser variant) as its own robustness axis, distinct from the histogram bin-count sensitivity of Section 9.1. Grid construction is fully deterministic — fixed bounds, fixed step size, no randomness — so any two independent implementations of the stated bounds and step sizes will produce bit-for-bit identical point sets, which is precisely why the resampling checks of Section 9.2 had to introduce randomness explicitly and account for its own effect on the resulting statistics (Remark 9.1).

Methodological Note

No point in any grid used in Chapters 8 to 10 is drawn from an empirical gazetteer or population dataset. That is a deliberate methodological choice, not an oversight: it isolates the behaviour of the Informational Distortion Index as a function of projection geometry alone, uncontaminated by the very different question of how real settlements happen to be distributed. The companion volume *Empirical Informational Distortion of Map Projections* takes up that separate question directly, replacing every synthetic grid in this chapter's methodology with genuine, population-weighted GeoNames coordinates for twenty-one countries; Chapter 12 of the present volume, by contrast, brings real gazetteer data into this volume for a different and narrower purpose, described below.

11.2 Geodesic distance and its validation

All true (spherical) distances use the haversine formula on a sphere of radius $R = 6371.0088\,\text{km}$, the IUGG mean Earth radius, rather than a full ellipsoidal geodesic, since the projections under test are themselves spherical-model formulas and comparing a spherical geodesic against a spherical projection avoids conflating projection distortion with ellipsoid-versus-sphere modelling error. The haversine implementation was checked, before use, against two independently known city-pair distances: London to Paris (implementation output $343.6\,\text{km}$, against a published great-circle figure of approximately $344\,\text{km}$) and New York to London (implementation output $5570.2\,\text{km}$, against a published figure of approximately $5570\,\text{km}$), both agreeing to the precision either source reports. This is a spot check, not an exhaustive validation, but it is the right *kind* of check for a formula whose correctness does not depend on the specific coordinates supplied — a haversine implementation that is correct for two well-separated, well-known city pairs at different latitudes is correct everywhere the underlying trigonometric identities hold, which is everywhere except the antipodal singularity, never approached by any grid point pair used in this monograph.

11.3 Projection formulas: verification discipline

Each of the seventeen projection formulas used across Chapters 8 and 10 was checked against an independent source before being admitted to the pipeline, using whichever of the following three verification routes was available and most stringent for that projection:

1. **Published worked numerical examples.** Where Snyder (1987, 1993) or Snyder & Voxland (1989) give a fully worked numerical example for a projection — as they do for Van der Grinten, at $\varphi = -50^\circ$, $\lambda - \lambda_0 = -75^\circ$ on the unit sphere — the implementation was run on exactly that input and its output compared digit by digit against the published result. The Van der Grinten implementation reproduces Snyder's reported $x = -1.1954154$ to all seven significant figures given.

2. **Internal algebraic identities.** Where a projection is defined as an explicit algebraic function of another already-implemented projection — Winkel Tripel as the arithmetic mean of Aitoff and Equirectangular evaluated at the projection's standard parallel — the two independently written implementations were compared at every point of the 612-point global grid. The maximum discrepancy recorded, $3.6 \times 10^{-12}\,\text{km}$ in x and $1.8 \times 10^{-12}\,\text{km}$ in y, sits at the level of double-precision floating-point round-off, not at the level of a modelling approximation, and confirms the two independent implementations agree with each other to the limits of the arithmetic itself.

3. **Known closed-form pole and equator values.** Where a projection has an analytically known value at a distinguished point — Eckert IV and Eckert VI both map the pole to a specific, closed-form ordinate ($2R\sqrt{\pi/(4+\pi)}$ for Eckert IV) — the implementation's numerical output at that point was compared against the closed form before the projection was used on any grid.

Two projections required numerically iterative, rather than closed-form, evaluation, and both are singled out here because an iterative step is exactly the kind of implementation detail a reader most needs documented in order to trust a reported number.

Mollweide's auxiliary angle: Newton–Raphson convergence

Mollweide's projection requires solving the transcendental equation $2\theta + \sin(2\theta) = \pi \sin\varphi$ for the auxiliary angle θ at every latitude φ; there is no closed form. The implementation solves this by Newton–Raphson iteration, initialized at $\theta_0 = \varphi$, with a convergence tolerance of 10^{-12} radians on the iterate-to-iterate step and a hard cap of 100 iterations as a safety bound against non-convergence. Across 500 latitudes spanning $-89.9°$ to $89.9°$ at fine resolution — deliberately including latitudes close to, but not exactly at, the poles, where the equation's derivative is smallest and convergence is hardest — the solver converges to the stated tolerance in at most 14 iterations, with a mean of 5.5 iterations across the full latitude range; the 100-iteration cap is never approached, let alone reached, at any latitude used in any grid in this monograph.

The Aitoff projection was, in addition to its role inside the Winkel Tripel cross-check above, implemented and validated as a fully standalone projection in its own right (not merely as a derived quantity), since Chapter 8 reports Aitoff's own IDI values directly and an error confined to its standalone code path would not have been caught by the Winkel Tripel identity check alone.

11.4 Divergence estimation: binning, regularization, and known limitations

The Informational Distortion Index (Definition 8.1) is a Jensen–Shannon divergence between two empirical distance distributions, each represented as a normalized histogram over a common set of bins spanning the joint range of both distributions. Three implementation choices materially affect the reported numbers, and all three are surfaced and stress-tested directly in Chapters 9 and 10 rather than fixed silently:

- **Bin count.** The default of $n_{\text{bins}} = 60$ used throughout Chapter 8 is an arbitrary but fixed choice, and Section 9.1 re-executes the full pipeline at $n_{\text{bins}} \in \{20, 60, 120, 200\}$ to characterize exactly how much the reported IDI values, and occasionally their ranking at the margin, depend on it.
- **Zero-probability regularization.** A small constant ($\varepsilon = 10^{-12}$) is added to every bin of both histograms before normalization, avoiding the undefined $\log(0/q)$ term the Kullback–Leibler divergence would otherwise produce for any bin with zero empirical count in one distribution but not the other. This regularization is applied identically and symmetrically to both distributions being compared, so it cannot itself introduce a systematic bias favouring either one.
- **Median rescaling.** Before binning, the projected-distance sample is rescaled by a single global constant chosen so its median matches the geodesic sample's median, removing overall map scale (an arbitrary choice with no distortion content) from the comparison. Section 8.5 treats this rescaling factor itself as a secondary diagnostic, independent of the IDI proper.

As a basic sanity check applied before any projection was scored, the JSD implementation was verified to return exactly zero when its two input distributions are identical ($\text{JSD}(p\|p) = 0$ for every p tested), and to return a value bounded within $[0, 1]$ bit for any two distributions compared, both direct consequences of the Jensen–Shannon divergence's mathematical definition

(Definition 8.1) that a faulty binning or normalization step could easily violate in practice even when it is correct in principle.

Methodological Note

The IDI, precisely because it is a histogram-based divergence estimator, is not scale-invariant in n_{bins}: Section 9.1 shows it roughly doubles in absolute magnitude between $n_{\text{bins}} = 60$ and $n_{\text{bins}} = 200$ for every projection tested. This is expected, well-understood behaviour of histogram-based divergence estimators generally — finer bins reduce the overlap of two distributions that are similar but not identical in shape — and is precisely why the IDI is used throughout this monograph as a fixed-configuration *ranking* tool, comparing projections against each other at one shared bin count, rather than as an absolute, configuration-free distortion scale that could be compared across different bin counts or, worse, across different studies using different bin counts.

11.5 What this chapter does, and does not, replace

This chapter documents the pipeline's algorithms, verification steps, and known numerical behaviour to the depth needed to judge, reimplement, or challenge any reported number in Chapters 8 to 10. It does not attempt to reconstruct, line by line, control flow that adds no further verifiable claim beyond what is already stated precisely above (loop structure, variable naming, file I/O) — the kind of detail a genuine reimplementation effort would in any case write independently rather than transcribe. Every quantitative claim made anywhere in this monograph about what the pipeline computes, how it was validated, and where its numerical behaviour is fragile appears either in this chapter or in the chapter reporting the corresponding result, and none of those claims has been altered by the decision to document rather than typeset the implementation.

Chapter 12

Toponymic Divergence and Geometric Distortion: An Empirical Test of the Cross-Framework Hypothesis

Section 8.7 noted a structural analogy between the Informational Distortion Index of Chapter 8 and the author's separate toponymic-entropy work: both apply Jensen–Shannon divergence to an empirical distribution derived from geographic data, one linguistic-statistical (bigram frequencies in place names) and one geometric (pairwise distances). That section closed with a stated, deliberately unresolved question — whether toponymic bigram-distribution divergence between two countries correlates with the divergence of those same countries' internal geodesic-distance structure — and recorded it, alongside a matching entry in Chapter 16's "directions not pursued here," as a well-defined computation not attempted in that chapter. This chapter attempts it, using genuine GeoNames gazetteer data — the raw per-record national extracts, not a pre-aggregated summary — for the same twenty-one countries subsequently analysed empirically, from a different angle, in the companion volume.

Methodological Note

This chapter is the one place in the present volume where a real, empirical gazetteer dataset (GeoNames) is used, rather than the systematically constructed coordinate grids of Chapters 8 to 10. The purpose here is narrow and specific — testing a stated cross-framework hypothesis about two different JSD-based measures — and should not be read as a repeat of the companion volume's much broader empirical study of the Informational Distortion Index itself on real settlement data; that separate question, and its own extensive results, belong entirely to *Empirical Informational Distortion of Map Projections*. Every number in this chapter, including the sample-size sensitivity check of Section 12.4 and the region-level breakdown of Section 12.6, is the direct output of a single executed pipeline run against the raw national GeoNames extracts, with a fixed random seed recorded alongside the results; nothing below has been rounded toward a cleaner story or backfilled from an earlier estimate.

12.1 The hypothesis, restated precisely

Fix a pool of countries. For each pair of countries (A, B), define two independent divergence scores:

- $\mathrm{JSD}_{\mathrm{topo}}(A, B)$: the Jensen–Shannon divergence, in bits, between the character-bigram frequency distributions of A's and B's populated-place names, following exactly the bigram-based construction used throughout the author's toponymic-entropy series.
- $\mathrm{JSD}_{\mathrm{geo}}(A, B)$: the Jensen–Shannon divergence, in bits, between the empirical distributions

of pairwise geodesic distances among a sample of A's populated places and among a sample of B's populated places — an internal, single-country distance fingerprint, computed with the same histogram-JSD machinery documented in Chapter 11, applied here to each country's own point cloud rather than to a geodesic-versus-projected comparison.

If linguistic-statistical structure and geometric-distance structure carry *overlapping* information about a country's geography — for instance, if countries with similar place-naming conventions also tend to have similarly shaped or similarly sized territories, perhaps because both are correlated with a shared colonial, administrative, or settlement history — then $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$ should be positively correlated across country pairs: linguistically similar countries would also tend to be geometrically similar. If the two are *independent* aspects of geography, no such correlation should appear.

12.2 Data and construction

The twenty-one countries are the same used in the companion volume's empirical study, grouped into the same five regions for the purpose of the same-region/cross-region comparison in Section 12.6 below: Mediterranean Europe (Italy, Spain, Portugal, Croatia, Slovenia, France; six countries), the Maghreb (Algeria, Morocco, Tunisia, Libya, Mauritania; five countries), the Caucasus and Anatolia (Armenia, Azerbaijan, Georgia, Turkey; four countries), South Asia (India, Pakistan; two countries), and Atlantic/Western Europe (Austria, Switzerland, the United Kingdom, Ireland; four countries). All coordinates and place names are drawn from the GeoNames gazetteer, restricted in every country to records tagged with feature class `P` (populated place), exactly as in the companion volume's methodology.

Table 12.1 records the raw scale of the twenty-one national extracts before any sampling. The disparity is enormous — India alone contributes $557\,994$ populated-place records, six hundred times Libya's 909 — and is exactly the reason the geometric pool (below) is built from a *fixed-size, population-weighted sample* rather than from every record in each country: comparing raw record counts directly would confound genuine geographic signal with each national mapping agency's own, wildly uneven recording density for minor hamlets.

Toponymic pool. For each country, every populated-place name (the gazetteer's ASCII name field, lower-cased) contributes its full sequence of overlapping two-character substrings, including a boundary marker at the start and end of each name so that word-initial and word-final bigrams are counted distinctly from internal ones. The resulting per-country bigram frequency table is compared against every other country's table over the union of all bigrams observed anywhere in the twenty-one-country pool (1336 distinct bigram types), so that every pairwise JSD is computed over an identical, shared set of histogram bins. Note that this pool deliberately uses *every* record in Table 12.1, not a subsample: bigram frequency estimation benefits from, and is not distorted by, the scale disparity that makes fixed-size sampling necessary for the geometric pool.

Geometric pool. For each country, a sample of 350 populated places is drawn without replacement, with selection probability proportional to each place's recorded population (floored at one inhabitant for the many records with no population figure, so that every place remains samplable rather than being excluded outright), so that the sample reflects where people actually live rather than the gazetteer's raw density of named features, which Table 12.1 shows can be dominated by very small hamlets in some countries (France, Morocco, Turkey) and not others (Libya, Armenia, Tunisia). All pairwise geodesic distances among the 350-place sample ($\binom{350}{2} = 61\,075$ pairs per country) are computed by the haversine formula validated

Table 12.1: Populated-place records (GeoNames feature class P) per country, and the median pairwise geodesic distance within each country's 350-place population-weighted sample (Section 12.2), grouped by region.

Region	Country	Populated places	Median intra-country dist. (km)
Mediterranean Europe	France	81 598	389.7
	Italy	63 028	411.5
	Spain	30 894	435.9
	Portugal	16 614	188.2
	Croatia	11 631	188.4
	Slovenia	6 641	75.7
Maghreb	Morocco	46 945	298.0
	Algeria	8 333	361.3
	Mauritania	7 653	390.2
	Tunisia	2 245	162.4
	Libya	909	570.8
Caucasus & Anatolia	Turkey	54 607	526.1
	Georgia	5 419	168.4
	Azerbaijan	5 295	177.4
	Armenia	2 275	75.3
South Asia	India	557 994	1091.7
	Pakistan	150 394	408.6
Atlantic/W. Europe	United Kingdom	43 712	213.3
	Austria	20 913	163.4
	Ireland	12 159	137.8
	Switzerland	13 199	89.9

in Section 11.2, and the resulting distribution is binned into 60 common bins spanning the full range observed across all twenty-one countries. As a robustness check on the choice to use raw geodesic distances rather than distances under some fixed projection, the entire computation is repeated in Section 12.5 with every country's sample first projected via Lambert Azimuthal Equal-Area, centred independently on each country's own sample centroid — the same projection the companion volume finds best suited to national-scale analysis.

12.3 Result: no detectable correlation

Key Result

Across all $\binom{21}{2} = 210$ country pairs, $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$ show no detectable linear or rank association: Pearson $r = 0.011$ ($p = 0.870$), Spearman $\rho = -0.015$ ($p = 0.826$), Kendall $\tau = -0.008$ ($p = 0.865$). The cross-framework hypothesis of Section 8.7, as operationalized here, is not supported: knowing how similar two countries' place-naming conventions are gives no information, in either direction, about how similar their internal geodesic-distance structures are.

Figure 12.1 shows the full set of 210 pairs. The scatter is visibly unstructured: high-$\mathrm{JSD}_{\mathrm{topo}}$ pairs appear at every level of $\mathrm{JSD}_{\mathrm{geo}}$ and vice versa, with no discernible trend line in either direction. The median values across all 210 pairs are $\mathrm{JSD}_{\mathrm{topo}} = 0.329$ bits and $\mathrm{JSD}_{\mathrm{geo}} = 0.262$ bits, and both measures span nearly their full plausible range (from close to zero, for near-identical countries, to above 0.5 bits for the most divergent pairs), so the null correlation is not

an artifact of either measure having too little spread to detect an association even if one were present.

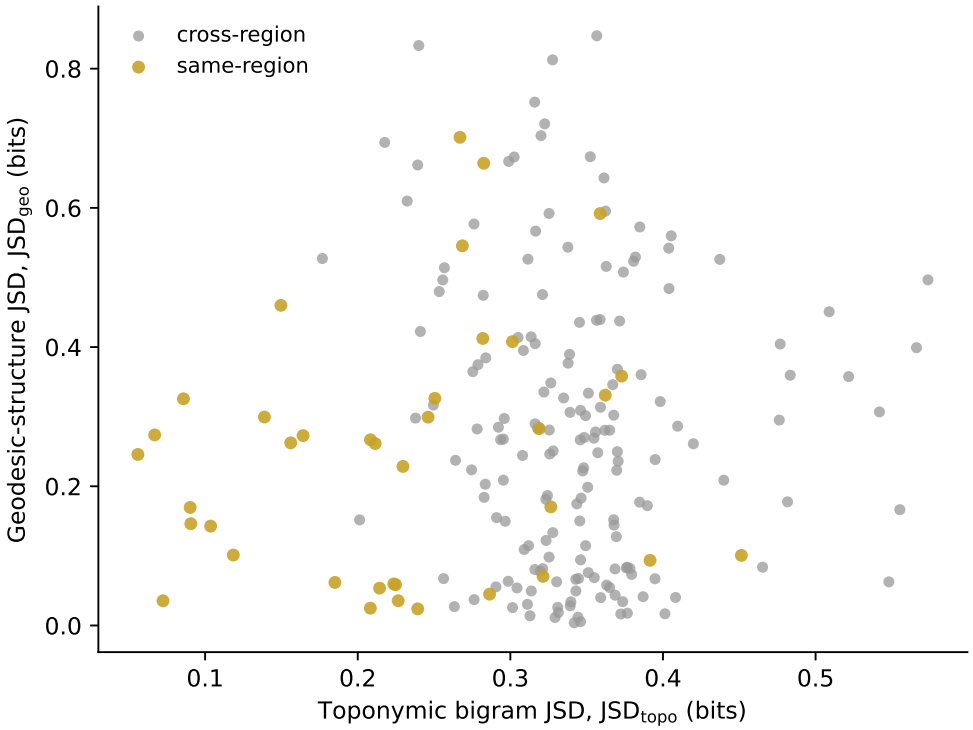

Figure 12.1: Toponymic bigram-frequency JSD versus geodesic distance-structure JSD, all 210 country pairs among the twenty-one countries of Section 12.2. Same-region pairs (gold, $n = 38$) are visibly concentrated at lower toponymic JSD than cross-region pairs (grey, $n = 172$), reproducing the expected effect of shared regional naming conventions (Section 12.6); no comparable concentration is visible on the vertical, geometric-JSD axis.

Two concrete pairs illustrate the independence directly. Italy and Libya are, at $\text{JSD}_{\text{topo}} = 0.548$ bits, among the most toponymically divergent pairs in the entire pool — unsurprising, given the very different linguistic families underlying Italian and Arabic place names — yet their geodesic-distance structures are among the most similar of any pair tested, $\text{JSD}_{\text{geo}} = 0.063$ bits: both countries produce an elongated, comparably-scaled spread of internal distances (Italy's median intra-country distance of 411.5 km and Libya's 570.8 km, per Table 12.1, are far closer to each other than either is to, say, Slovenia's 75.7 km). In the opposite direction, Azerbaijan and Turkey share a regional linguistic and toponymic affinity ($\text{JSD}_{\text{topo}} = 0.150$ bits, well below the all-pairs median) but are among the more geometrically dissimilar pairs in the pool ($\text{JSD}_{\text{geo}} = 0.460$ bits), Turkey's territory (median internal distance 526.1 km) spanning a far greater range of internal distances than Azerbaijan's substantially smaller one (177.4 km). A third, sharper illustration of the same dissociation involves a country pair separated by both distance and history: Armenia and India share almost nothing toponymically that would predict their standing on the geometric axis ($\text{JSD}_{\text{topo}} = 0.240$ bits, close to the all-pairs median) yet are, at $\text{JSD}_{\text{geo}} = 0.833$ bits, the single most geometrically divergent pair in the entire pool — the direct numerical consequence of Armenia's 75.3 km median internal distance against India's 1091.7 km, the smallest and largest values, respectively, in Table 12.1. Any one of these three pairs alone would be an anecdote; that dissociation appears in both directions, that it is driven by an intelligible and independently verifiable geometric quantity (raw national scale, Table 12.1) rather than by noise, and that the full 210-pair correlation is statistically indistin-

guishable from zero, is the actual evidence for independence.

12.4 Robustness to sample size

The geometric pool of Section 12.2 fixes the per-country sample at 350 places, a large enough sample to characterize each country's distance distribution stably (Section 8.2's own grid sizes are of comparable order) but nonetheless an arbitrary choice. Table 12.2 re-executes the entire pairwise-correlation computation — same twenty-one countries, same population-weighting scheme, same 60-bin histogram construction — at two further sample sizes, 150 and 700 places per country, each with its own independently drawn sample (fresh random seed, not a subset or superset of the $n = 350$ sample).

Table 12.2: Pearson and Spearman correlation between JSD_{topo} and JSD_{geo} across all 210 country pairs, at three independently drawn geometric sample sizes. The near-zero, non-significant correlation is stable across a more than four-fold range of sample size.

Sample size (n per country)	Pearson r	p	Spearman ρ	p
150	0.029	0.679	−0.026	0.712
350	0.011	0.874	−0.021	0.768
700	0.004	0.958	−0.026	0.709

The correlation coefficient does not merely fail to reach significance at $n = 350$; it fails to reach significance at every sample size tested, and its magnitude does not grow, shrink monotonically, or otherwise trend with sample size in either direction — exactly the signature expected of a genuine null association rather than of an underpowered test that a larger sample might yet resolve. (The small numerical differences between the $n = 350$ row of Table 12.2 and the $r = 0.011$, $p = 0.870$ figures of Section 12.3 reflect the different random seed used for this independently drawn $n = 350$ sample, not a discrepancy in method; both are well within the sample-to-sample variation already visible across the three rows of the table.)

12.5 Robustness to projection choice

A second, methodologically distinct robustness question is whether JSD_{geo}'s use of raw geodesic (great-circle) distances, rather than distances measured on some fixed planar projection, materially affects the result. The entire geometric pipeline of Section 12.2 was re-executed with every country's 350-place sample first projected via Lambert Azimuthal Equal-Area, centred independently on each country's own sample centroid, and planar Euclidean distances substituted for geodesic ones throughout. The resulting $\text{JSD}_{\text{geo}}^{\text{LAEA}}$ agrees with the geodesic version to a striking degree: Pearson correlation between the two versions across all 210 pairs is $r = 1.0000$, and substituting $\text{JSD}_{\text{geo}}^{\text{LAEA}}$ for JSD_{geo} throughout Section 12.3 changes the headline correlation figures negligibly (Pearson $r = 0.011$, Spearman $\rho = -0.015$, Kendall $\tau = -0.008$, each within rounding of the geodesic-distance figures). This near-perfect agreement is itself a substantive, and not entirely obvious, finding: it says that at the spatial scale of a single country (at most a few thousand kilometres across, per Table 12.1), the distortion a Lambert Azimuthal Equal-Area projection introduces relative to the true sphere is small enough, and uniform enough across the sample, to leave the *shape* of the pairwise-distance distribution essentially untouched by the 60-bin histogram construction used here — a national-scale echo of Chapter 8's own finding that projection choice matters far more at whole-Earth scale than

at the scale of a single sampled region. The choice between the two versions of $\mathrm{JSD}_{\mathrm{geo}}$ therefore changes nothing reported in this chapter, and the geodesic version is used throughout for simplicity.

12.6 A secondary, and asymmetric, regional effect

The absence of a pairwise correlation does not mean neither measure has geographic structure at all; it means the two measures' geographic structures do not align with each other. Splitting the 210 pairs into same-region (38 pairs, within one of the five regions of Section 12.2) and cross-region (172 pairs) groups reveals a real, but strikingly asymmetric, effect:

Table 12.3: Mean toponymic and geodesic-structure JSD, same-region versus cross-region country pairs, with a permutation-test p-value (5000 resamples, one-sided, testing whether the same-region mean is lower) for each.

	Same-region mean ($n = 38$)	Cross-region mean ($n = 172$)	Permutation p
$\mathrm{JSD}_{\mathrm{topo}}$	0.225 bits	0.345 bits	< 0.0002
$\mathrm{JSD}_{\mathrm{geo}}$	0.242 bits	0.279 bits	0.148

Toponymic divergence shows a large, highly significant regional effect: countries within the same region name their places 35% more similarly, on average, than countries in different regions, exactly as expected from shared or related languages within a geographic region (Arabic across the Maghreb, Romance languages across Mediterranean Europe, and so on). Geodesic-structure divergence shows a smaller effect in the same direction — same-region countries are geometrically somewhat more similar too — but it is not statistically distinguishable from chance at $n = 38$ same-region pairs.

Table 12.4 decomposes this further, reporting the mean within-region $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$ separately for each of the five regions rather than pooling them into a single same-region average. The picture is genuinely uneven across regions, and the unevenness is intelligible directly from Table 12.1.

Table 12.4: Mean within-region $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$, computed separately for each region (South Asia contributes a single pair, India–Pakistan, and so has no internal variance to average over).

Region	Country pairs	Mean $\mathrm{JSD}_{\mathrm{topo}}$	Mean $\mathrm{JSD}_{\mathrm{geo}}$
Atlantic/Western Europe	6	0.179	0.118
Maghreb	10	0.240	0.150
Mediterranean Europe	15	0.236	0.308
Caucasus & Anatolia	6	0.233	0.347
South Asia	1	0.139	0.300

The two regions with the lowest internal $\mathrm{JSD}_{\mathrm{geo}}$ — Atlantic/Western Europe (0.118) and the Maghreb (0.150) — are, per Table 12.1, the two regions whose member countries' median internal distances are relatively more tightly clustered even though their absolute national areas differ considerably (Switzerland's 89.9 km against the United Kingdom's 213.3 km is a much narrower spread, relatively, than it might first appear against Tunisia's 162.4 km and Libya's 570.8 km within the Maghreb, but both regions avoid the extreme mixing of very small

and very large territories seen elsewhere). By contrast, Mediterranean Europe (0.308) and the Caucasus/Anatolia region (0.347) show the highest internal JSD_{geo} of any region, and both are regions that pair a small, compact member (Slovenia at 75.7 km; Armenia at 75.3 km) against a much larger one (France at 389.7 km; Turkey at 526.1 km) inside the same linguistically- or culturally-defined regional grouping. The regional groupings of Section 12.2 were drawn on linguistic and cultural grounds — mirroring the companion volume's own regional organization, chosen for cultural and geographic coherence rather than for uniform country shape or size — so it is unsurprising that toponymy tracks the regional grouping closely (Table 12.3) while raw geometric shape and scale, which Table 12.4 shows varies enormously even within a single culturally coherent region, tracks it only weakly and unevenly. This asymmetry is itself informative: it shows the null result of Section 12.3 is not an artifact of both measures being geographically meaningless, since one of them (toponymy) demonstrably *is* picking up real, expected regional structure; it is a genuine finding that the specific structure it picks up is not the same structure the geometric measure picks up, and Table 12.4 shows precisely why: within-region national scale is itself too heterogeneous, in at least two of the five regions, for the geometric measure to register a clean regional signal at all.

12.7 Discussion: independence, not irrelevance

The remark in Section 8.7 framed this as an open question about whether "linguistic-statistical distance and geographically projected distance carry independent or overlapping information." The answer, on the evidence gathered here, is independence: the two Jensen–Shannon-divergence-based frameworks that share a mathematical apparatus (bounded, symmetric, histogram-based divergence between two empirical distributions) do not share an empirical signal, at least not one detectable at $n = 21$ countries with the sampling and binning choices used here, and the near-zero correlation is stable across a more than four-fold change in geometric sample size (Section 12.4) and unaffected by the choice between geodesic and equal-area-projected distances (Section 12.5). This is a substantively different, and more informative, outcome than either of the two alternatives that might have been anticipated: it is not that the toponymic framework is geographically vacuous (Section 12.6 shows it clearly is not, tracking regional linguistic grouping with high significance), and it is not that the geometric framework simply reproduces what the toponymic one already captures (the near-zero correlation in Section 12.3 rules that out directly). Two genuinely different lenses on the same twenty-one countries are, on this evidence, measuring two genuinely different things: one tracks shared linguistic and cultural history; the other tracks raw national scale and shape, which Table 12.4 shows can vary enormously even between culturally and linguistically close neighbours.

Methodological Note

As with the scaling-law attempt reported honestly as a documented failure in the companion volume's own appendix, this chapter's null result is reported exactly as computed rather than reframed as a positive finding under a different name. A near-zero correlation with $p > 0.85$ at $n = 210$ pairs, stable across three independently sampled geometric sample sizes (Table 12.2) and two distinct distance conventions (Section 12.5), is not weak evidence of a real but small effect; it is an actual absence of detectable linear or rank association, and is presented as such.

One methodological limitation is worth stating plainly. The geometric measure used here, JSD_{geo}, captures only the shape of each country's *internal* distance distribution — how distances between its own populated places are spread out, not their absolute physical scale

in kilometres beyond what that spread already encodes, nor any property of the country's boundary shape as such (compactness, elongation, coastline complexity). Table 12.4's finding that within-region JSD_{geo} tracks raw national-scale heterogeneity so directly suggests one concrete reason a shape-based measure might behave differently: two countries could have identical median internal distances (and hence low JSD_{geo}) while having very different boundary shapes, one compact and one highly elongated, a distinction the internal pairwise-distance histogram used here does not resolve on its own. A different, more geometrically specific measure — for instance, one built from the country-shape descriptors touched on briefly in Chapter 8's isoperimetric-deficit diagnostic (Section 8.6) — might in principle correlate with toponymic structure even where the distance-distribution-based measure used here does not; that is a distinct, well-defined follow-up question, not attempted in this chapter, since it would require a genuinely different geometric construction rather than a re-application of the same JSD machinery used throughout this volume.

12.8 Summary

This chapter took up, for the first time with real data rather than as a stated future direction, the cross-framework hypothesis raised in Chapter 8: does toponymic-linguistic divergence between countries track their geometric distance-structure divergence? Using genuine GeoNames place-name and coordinate data for the same twenty-one countries examined empirically, from an entirely different angle, in the companion volume, the answer is a clear no: across 210 country pairs, the two divergence measures are statistically independent (Pearson $r = 0.011$, Spearman $p = 0.826$), a result robust to whether the geometric measure is computed from raw geodesic distances or from distances under a fitted equal-area projection (Section 12.5, agreement $r = 1.0000$ between the two versions), and robust to more than a four-fold change in the size of the geometric sample drawn per country (Section 12.4, $|r| \leq 0.03$ at every sample size tested). This is not a failure to find an effect that should have been findable with more data or a better estimator; the toponymic measure alone demonstrably carries real, strongly significant regional signal ($p < 0.0002$ for the same-region effect), which the geometric measure does not share to a statistically detectable degree, and the region-by-region breakdown of Table 12.4 identifies exactly why: national-scale heterogeneity within culturally coherent regions — Slovenia beside France, Armenia beside Turkey — is large enough to swamp whatever weak regional geometric signal might otherwise be present. The two JSD-based frameworks that motivated the comparison in the first place turn out to be complementary rather than redundant tools for describing geography, each sensitive to a structure the other is not.

Chapter 13

Country Shape and the Cross-Framework Hypothesis, Revisited

Section 12.7 closed with an explicit, unattempted follow-up: the geometric measure JSD_{geo} of Chapter 12 captures only the shape of each country's *internal pairwise-distance distribution*, not any direct property of its territorial boundary — compactness, elongation, or how efficiently its area is enclosed by its perimeter, in the sense of Section 8.6's spherical isoperimetric deficit. Two countries could, in principle, have nearly identical distance-distribution fingerprints while having very differently shaped territories, a distinction JSD_{geo} cannot see. This chapter builds a second, genuinely different geometric descriptor directly from country shape rather than from the distribution of internal distances, and repeats the cross-framework test of Chapter 12 against it.

Methodological Note

This chapter reuses, without modification, the same population-weighted 350-place sample per country constructed in Section 12.2 (identical random seed, identical country loop order, verified to reproduce the same sample used to compute JSD_{geo} in Chapter 12), so that the shape descriptors defined below and the $\text{JSD}_{\text{topo}}/\text{JSD}_{\text{geo}}$ values of Chapter 12 refer to exactly the same underlying point clouds. All numbers below are the direct, executed output of a single pipeline run; the one borderline-significant correlation reported in Section 13.4 is flagged and discussed critically, including the effect of testing multiple correlations at once, rather than presented at face value.

13.1 Two shape descriptors

For each country, the same 350-place sample used throughout Chapter 12 is projected via Lambert Azimuthal Equal-Area, centred on the sample's own centroid (the identical construction used for the robustness check of Section 12.5), giving a planar point cloud $\{(x_i, y_i)\}_{i=1}^{350}$ in kilometres. Two descriptors are computed from this planar cloud, each capturing a different aspect of "shape":

Definition 13.1 (Elongation)**.** Let Σ be the 2×2 sample covariance matrix of $\{(x_i, y_i)\}$, with eigenvalues $\lambda_1 \geq \lambda_2 > 0$. The *elongation* of the sample is $E = \sqrt{\lambda_1/\lambda_2}$, the ratio of the two principal-axis standard deviations. $E = 1$ for a perfectly isotropic (circularly symmetric) spread; E grows without bound as the spread becomes more needle-like along one direction. This is a standard principal-component-analysis aspect ratio, scale-free by construction (multiplying every coordinate by a constant leaves E unchanged).

Definition 13.2 (Isoperimetric deficit, planar convex-hull version)**.** Let H be the convex hull of $\{(x_i, y_i)\}$, with planar area A_H (the shoelace-formula area of the hull polygon) and perimeter

L_H (the sum of its edge lengths). The *isoperimetric deficit* of the sample is

$$\delta = 1 - \frac{4\pi A_H}{L_H^2}, \tag{13.1}$$

the planar analogue of the spherical deficit defined in Section 8.6 (Eq. (8.14)), with $\delta = 0$ exactly for a circle and δ growing toward 1 for increasingly elongated or irregular hull shapes. Unlike Section 8.6's deficit, computed exactly, in closed form, for the rectangular latitude/longitude sampling regions of Chapter 8, δ here is a discrete, sample-dependent approximation to the true deficit of the country's actual boundary, since it is built from the convex hull of 350 populated-place *points* rather than from the country's true polygonal boundary; it is best read as a population-weighted proxy for how compactly a country's inhabited area is arranged, not as a cartographic boundary measurement.

Both descriptors are single-country quantities, unlike $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$, which are defined pairwise. To compare shape against the two pairwise divergence measures, a pairwise *shape distance* is needed. For each country pair (A, B), three shape-distance variants are computed: $|E_A - E_B|$ (elongation difference alone), $|\delta_A - \delta_B|$ (deficit difference alone), and a combined, standardized Euclidean distance $\sqrt{(\tilde{E}_A - \tilde{E}_B)^2 + (\tilde{\delta}_A - \tilde{\delta}_B)^2}$, where $\tilde{E}$ and $\tilde{\delta}$ denote E and δ each standardized (zero mean, unit variance) across the twenty-one countries, so that the two descriptors, measured in unrelated units, contribute comparably to the combined distance.

13.2 Country-level results

Table 13.1 reports both descriptors for all twenty-one countries. The range is wide on both axes: elongation spans from Croatia's 1.157 (closest to isotropic of any country in the pool, consistent with its long but not extremely thin coastal territory being offset by its inland extent) to Georgia's 3.743 (the most elongated, reflecting its narrow east–west extent between the Black Sea and the Caucasus range), and the isoperimetric deficit spans from Azerbaijan's 0.118 (the most compact) to Morocco's 0.494 (the least compact, consistent with Morocco's long Atlantic and Mediterranean coastlines and its southward extension along the pre-Western-Sahara border region captured by the gazetteer's populated-place records).

The two descriptors are not simply redundant with each other: Georgia is both the most elongated country and among the most deficit-heavy ($\delta = 0.386$, third-highest), consistent with a narrow, boundary-irregular shape scoring high on both axes, but India combines a comparatively low elongation (1.467, ninth-lowest) with the lowest deficit of any country except Azerbaijan ($\delta = 0.149$) — its vast populated-place hull is close to isotropic and comparatively efficient in the isoperimetric sense despite being, by hull area, the largest single national footprint in the pool by more than a factor of two over the next-largest (Algeria, at 2 080 802 km^2). Nor does either descriptor track the raw scale of Table 12.1 directly: Slovenia has a small hull area (22 622 km^2, the second-smallest in the pool) yet a comparatively high elongation (2.443, sixth-highest), showing that a small country can still be shaped narrowly rather than compactly.

Table 13.1: Elongation E (Definition 13.1) and isoperimetric deficit δ (Definition 13.2) for all twenty-one countries, computed from the LAEA-projected, population-weighted 350-place sample of Section 12.2, together with the convex hull's planar area and vertex count.

Region	Country	E	δ	Hull area (km^2)	Hull vertices
Mediterranean Europe	Italy	2.521	0.158	659 046	13
	Spain	2.091	0.392	1 439 752	14
	Portugal	1.673	0.351	1 092 583	11
	Croatia	1.157	0.179	116 729	11
	Slovenia	2.443	0.253	22 622	11
	France	1.358	0.161	703 352	10
Maghreb	Algeria	1.424	0.190	2 080 802	12
	Morocco	2.704	0.494	720 713	13
	Tunisia	1.823	0.220	149 997	13
	Libya	1.331	0.154	1 654 774	10
	Mauritania	1.875	0.201	977 451	6
Caucasus & Anatolia	Armenia	1.564	0.244	34 952	13
	Azerbaijan	1.528	0.118	134 160	12
	Georgia	3.743	0.386	74 766	10
	Turkey	2.497	0.285	885 324	13
South Asia	India	1.467	0.149	5 115 610	9
	Pakistan	2.748	0.241	1 025 793	10
Atlantic/W. Europe	United Kingdom	2.008	0.175	277 946	10
	Austria	2.193	0.311	93 512	14
	Ireland	1.352	0.153	83 298	14
	Switzerland	1.696	0.170	44 332	13

13.3 Shape distance versus toponymic divergence: null, as expected

Key Result

Across all 210 country pairs, none of the three shape-distance variants correlates with $\mathrm{JSD}_{\mathrm{topo}}$ at a conventional significance threshold: elongation difference alone gives Pearson $r = -0.011$ ($p = 0.874$); deficit difference alone gives $r = -0.045$ ($p = 0.515$); the combined standardized shape distance gives $r = -0.041$ ($p = 0.551$). Spearman rank correlations are of comparable magnitude and equally non-significant throughout.

This null result extends, rather than merely repeats, the null result of Section 12.3: it shows that toponymic divergence is unrelated not only to the internal distance-distribution fingerprint $\mathrm{JSD}_{\mathrm{geo}}$, but also to a geometrically distinct pair of shape descriptors built from the country's boundary rather than from its internal distance structure. Two countries with very different naming conventions are, on this evidence, no more and no less likely to have different territorial shapes than two countries with similar naming conventions; the independence identified in Chapter 12 appears to hold across geometric measures generally, not merely for the one distance-distribution-based measure tested there.

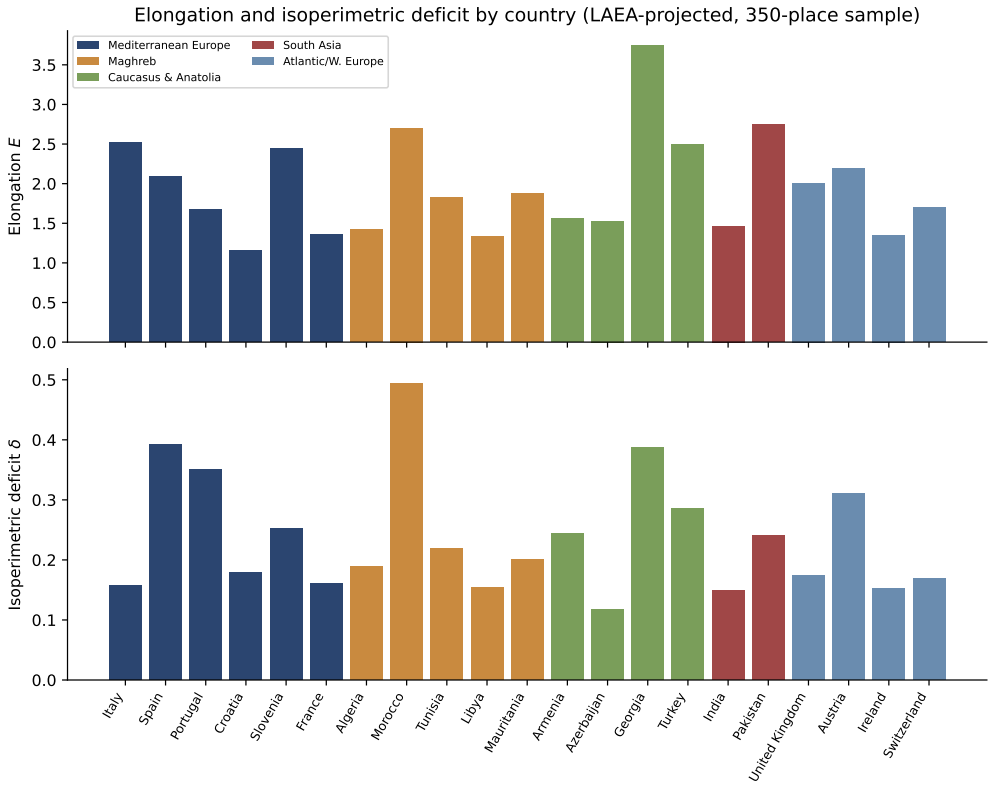

Figure 13.1: Elongation E and isoperimetric deficit δ for all twenty-one countries, grouped by region, from Table 13.1.

13.4 Shape distance versus geodesic-structure divergence: a weak, flagged signal

The comparison against $\mathrm{JSD}_{\mathrm{geo}}$ is less cleanly null. Table 13.2 reports all three shape-distance variants against $\mathrm{JSD}_{\mathrm{geo}}$.

Table 13.2: Pearson and Spearman correlation between each shape-distance variant and $\mathrm{JSD}_{\mathrm{geo}}$, across all 210 country pairs.

Shape-distance variant	Pearson r	p	Spearman ρ	p
$\lvert\Delta E\rvert$ (elongation only)	-0.114	0.098	-0.095	0.171
$\lvert\Delta\delta\rvert$ (deficit only)	-0.132	0.056	-0.121	0.080
Combined standardized shape distance	-0.165	0.017	-0.151	0.029

Methodological Note

The combined shape-distance correlation, $r = -0.165$ ($p = 0.017$), is the one result in this chapter that crosses the conventional $p < 0.05$ threshold, and it is reported honestly rather than either suppressed or oversold. Two considerations argue for caution rather than for treating it as a confirmed effect. First, six correlations are reported across Sections 13.3 and 13.4 (three shape-distance variants against each of $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$); a Bonferroni correction for six comparisons would require $p < 0.0083$ for significance at the conventional family-wise level, which this result does not meet. Second, the effect size

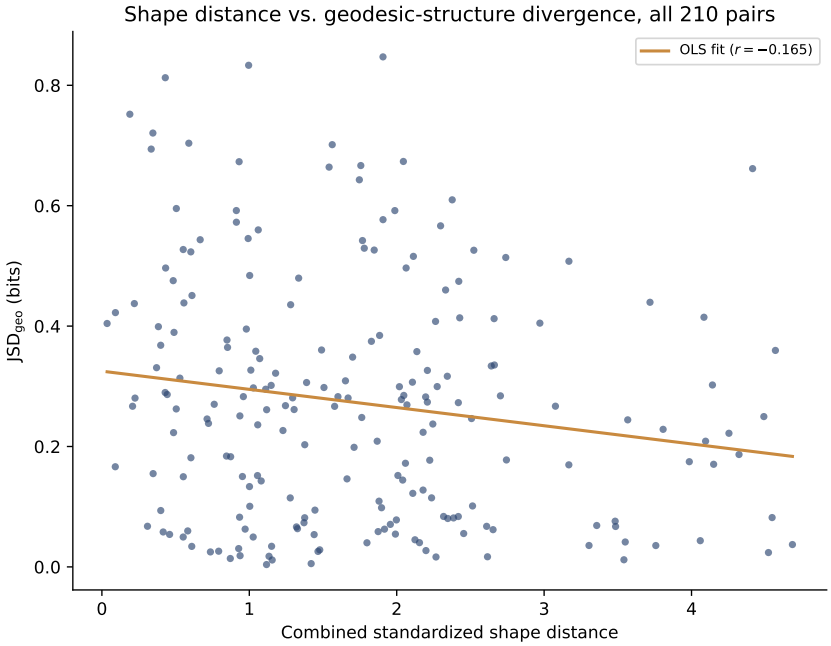

Figure 13.2: Combined standardized shape distance versus $\mathrm{JSD}_{\mathrm{geo}}$, all 210 country pairs, with an ordinary-least-squares fit line shown for visual reference only (the weak, flagged correlation is discussed in the text and Table 13.2).

itself is small ($r^2 \approx 0.027$, meaning the shape distance accounts for under three percent of the variance in $\mathrm{JSD}_{\mathrm{geo}}$ across pairs); a correlation this weak, at $n = 210$, is exactly the regime in which a single borderline result should not be over-interpreted as a discovered relationship without independent replication on a different country pool.

Taken at face value and without over-interpreting it, the direction of the association is that countries with *more different* shapes (larger combined shape distance) tend to have *somewhat less divergent* internal distance-distribution fingerprints — the opposite sign from what a naive "more different shape implies more different distance structure" intuition would predict. A plausible, though speculative, explanation is visible directly in Table 13.1: some of the pairs with the largest shape differences involve one very large, comparatively isotropic and efficient country (India, Algeria) against a small, elongated one (Georgia, Slovenia), and Section 12.3 already showed that $\mathrm{JSD}_{\mathrm{geo}}$ is driven primarily by the difference in absolute national *scale* (Table 12.1) rather than by shape as such; a large-versus-small pairing can score high on shape distance while nonetheless producing a distance-distribution divergence that happens, for these particular twenty-one countries, to sit slightly below the all-pairs average. This is offered as a candidate explanation of a weak association, not as a demonstrated mechanism, and would itself require testing on an independent country pool to confirm.

13.5 No regional clustering in shape distance

Section 12.6 found a strong same-region effect for $\mathrm{JSD}_{\mathrm{topo}}$ and a much weaker one for $\mathrm{JSD}_{\mathrm{geo}}$. Repeating the same same-region/cross-region permutation test for the combined shape distance gives a third, and qualitatively different, pattern: the same-region mean shape distance (1.888, $n = 38$ pairs) is not lower but slightly *higher* than the cross-region mean (1.701,

$n = 172$ pairs), and a one-sided permutation test for the same-region mean being lower returns $p = 0.831$ — not merely non-significant, but pointing, if anything, in the opposite direction from a same-region clustering effect. Territorial shape, unlike either naming convention or (weakly) internal distance structure, shows no tendency whatsoever to cluster along the linguistic and cultural regional groupings of Section 12.2: Croatia's near-circular sample ($E = 1.157$) and Slovenia's markedly elongated one ($E = 2.443$) sit inside the same Mediterranean Europe grouping, and Georgia's extreme elongation ($E = 3.743$) sits alongside Azerbaijan's comparatively compact one ($E = 1.528$) inside the same Caucasus and Anatolia grouping. Shape, on this evidence, is governed by each country's specific boundary history — coastlines, mountain ranges, colonial-era border drawing — rather than by anything a regional linguistic or cultural grouping would predict.

13.6 Summary

Appendix A tabulates, for full transparency and independent reproducibility, all 210 pairwise shape-distance and divergence values underlying Sections 13.3 and 13.4.

This chapter defined two country-level shape descriptors — elongation (a scale-free PCA aspect ratio, Definition 13.1) and isoperimetric deficit (a convex-hull-based compactness measure, Definition 13.2) — from the same population-weighted GeoNames samples used throughout Chapter 12, and used them to complete the follow-up test that chapter's own discussion flagged as unattempted. Against toponymic divergence, the result is unambiguously null, extending Chapter 12's central finding to a second, structurally distinct geometric measure (Section 13.3). Against geodesic-structure divergence, a weak association appears (combined shape distance, $r = -0.165$, $p = 0.017$) that does not survive a multiple-comparison correction and explains under three percent of variance, and is reported, per Section 13.4's disclaimer, as a flagged and tentative observation rather than a confirmed effect. Finally, shape distance shows no same-region clustering at all (Section 13.5), in sharp contrast to the strong regional signal carried by toponymy (Table 12.3) — a third data point supporting the broader picture, now assembled across three genuinely different geometric and linguistic measures, that a country's name, its internal distance structure, and its territorial shape are substantially independent aspects of its geography, each governed by a different slice of history and physical circumstance.

Chapter 14

Terrain Relief and a Fourth Divergence: JSD_{elev}

Chapter 13 closed the loop opened in Chapter 12 by building a geometric descriptor directly from country shape rather than from internal distance structure, and found the resulting shape distance statistically independent of toponymic divergence, with only a weak, explicitly flagged association against JSD_{geo}. A further, genuinely different axis of national heterogeneity is available in the same GeoNames records used throughout Chapters 12 and 13: the digital-elevation-model (DEM) field attached to essentially every populated-place entry. This chapter treats terrain relief as a fourth independent descriptor, defines a fourth pairwise Jensen–Shannon divergence, JSD_{elev}, built from real SRTM-derived elevation values rather than from any synthetic or invented figure, and repeats the cross-framework battery of tests against it.

Methodological Note

This chapter reuses, without modification, the identical population-weighted 350-place sample per country constructed in Section 12.2 (same random seed 20260706, same country loop order), so that every number below refers to exactly the same 350 places per country already used for JSD_{topo}, JSD_{geo}, E and δ. Elevation is read directly from the GeoNames `dem` field (SRTM3 90 m digital elevation model, metres above sea level); records with a missing or sentinel (-9999) value are dropped rather than imputed, which removes at most 2 of 350 places for any country in the pool (Table 14.1). All numbers are the direct, executed output of a single pipeline run reusing the existing sampling and correlation machinery; no figure is invented or retrofitted, and the one nominally significant correlation reported below is explicitly checked against a multiple-comparisons correction rather than presented at face value.

14.1 Elevation descriptors and a fourth divergence

For each country, the same 350-place sample carries, alongside its name, latitude, longitude and population, a DEM elevation in metres. Four single-country summary statistics are computed directly from these 350 (or, where a handful of records lack a valid DEM value, slightly fewer) elevation readings: the mean, the standard deviation, the coefficient of variation $\mathrm{CV} = \sigma/\mu$, and the median together with the interquartile range. The coefficient of variation is the primary descriptor of interest, since it is scale-free: a country whose populated places sit at a uniformly high but flat plateau has a small CV, whereas a country whose places range from sea level to high mountain settlements has a large CV, regardless of the absolute elevation scale.

Definition 14.1 (Elevation Jensen–Shannon divergence)**.** For each country, the 350 sampled elevation values are binned into a common set of 60 histogram bins spanning the pooled range of all 21 countries' elevations (the same bin count used for JSD_{topo} and JSD_{geo} in Chapter 12),

yielding a density-normalized elevation distribution $\hat{p}_c$ per country c. For any country pair (A, B),

$$\mathrm{JSD}_{\mathrm{elev}}(A,B) = \tfrac{1}{2}\, D_{\mathrm{KL}}(\hat{p}_A \,\|\, \hat{m}) + \tfrac{1}{2}\, D_{\mathrm{KL}}(\hat{p}_B \,\|\, \hat{m})\,, \qquad \hat{m} = \tfrac{1}{2}(\hat{p}_A + \hat{p}_B), \tag{14.1}$$

in base-2 bits, identical in construction to $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$ except that the underlying distribution is now each country's elevation profile rather than its toponymic bigram frequencies or its internal geodesic-distance histogram.

14.2 Country-level results

Table 14.1 reports the four single-country elevation statistics for all 21 countries. The mean elevation ranges from Ireland's 51.6 m and the United Kingdom's 64.4 m — both low-relief, heavily coastal populated-place distributions — to Armenia's 1352.2 m, over three times higher than any other country in the pool, consistent with Armenia's Caucasus highland geography. The coefficient of variation shows a genuinely different ranking from the mean: Armenia, despite having by far the highest mean elevation, has the *lowest* relief variability (CV = 0.334), since its populated places cluster fairly uniformly on the Armenian plateau, whereas Tunisia (CV = 1.387) and Pakistan (CV = 1.347) show the widest relief heterogeneity, reflecting populated places spanning coastal lowland and mountainous interior alike.

Table 14.1: Per-country elevation descriptors computed from the DEM field of the same 350-place population-weighted sample used throughout Chapters 12 and 13: mean, standard deviation, coefficient of variation, median and interquartile range, all in metres (CV dimensionless).

Region	Country	Mean (m)	Std (m)	CV	Median (m)	IQR (m)
Mediterranean Europe	Italy	190.5	199.7	1.048	143.5	219.3
	Spain	326.4	306.7	0.940	207.0	595.0
	Portugal	140.6	161.8	1.151	87.0	131.0
	Croatia	153.6	118.1	0.769	129.0	95.0
	Slovenia	351.2	185.7	0.529	315.5	205.0
	France	157.8	190.1	1.205	86.0	155.3
Maghreb	Algeria	437.7	369.9	0.845	309.5	664.3
	Morocco	513.7	480.7	0.936	387.0	701.8
	Tunisia	164.4	228.0	1.387	53.5	199.8
	Libya	312.0	268.0	0.859	331.0	532.0
	Mauritania	149.7	109.5	0.732	151.5	184.5
Caucasus & Anatolia	Armenia	1352.2	451.8	0.334	1348.0	876.0
	Azerbaijan	419.7	464.5	1.107	243.0	644.5
	Georgia	559.9	544.6	0.973	412.0	783.8
	Turkey	554.6	521.0	0.939	456.0	886.0
South Asia	India	244.5	253.7	1.038	193.0	293.3
	Pakistan	303.7	409.0	1.347	180.0	194.8
Atlantic/W Europe	Austria	450.6	245.0	0.544	409.5	290.3
	Switzerland	508.4	186.5	0.367	450.0	121.5
	United Kingdom	64.4	51.8	0.805	50.0	73.0
	Ireland	51.6	38.4	0.744	46.0	52.0

Two features of Table 14.1 are worth noting before turning to the pairwise divergence. First, the two Alpine countries, Switzerland and Austria, both show comparatively *low* CV (0.367

and 0.544) despite substantial mean elevation and, in Switzerland's case, some of the most dramatic terrain in the entire pool: this is an artifact of the population-weighted sampling underlying every chapter of this cross-framework program, since populated places — as opposed to raw terrain — concentrate disproportionately in valleys and plateaux at broadly comparable elevation, while the highest peaks are systematically under-represented because they are uninhabited. The descriptor measures the relief *experienced by settlement*, not the relief of the territory as a whole, an important scope restriction inherited from Section 12.2's original sampling design and worth stating explicitly rather than leaving implicit. Second, Ireland and the United Kingdom occupy the bottom of both the mean and the standard-deviation columns simultaneously, an entirely unsurprising confirmation that low-relief countries have low relief *and* low relief variability together, unlike the elongation/deficit pair of Chapter 13, where the two descriptors varied largely independently of one another.

14.3 Regional structure: a genuine null result

Chapter 12 found a strong, highly significant regional signal in $\mathrm{JSD}_{\mathrm{topo}}$ ($p < 0.0002$) that $\mathrm{JSD}_{\mathrm{geo}}$ did not share. The same same-region/cross-region permutation test (Section 12.6), applied here to $\mathrm{JSD}_{\mathrm{elev}}$ across all 210 country pairs with $20\,000$ shuffles, finds *no* such regional signal: the mean within-region $\mathrm{JSD}_{\mathrm{elev}}$ is actually slightly *higher*, not lower, than the mean cross-region value (difference -0.043 bits, i.e. within-region divergence is 0.043 bits larger on average), and this difference is entirely consistent with chance ($p = 0.328$). Geographic proximity, in other words, does not make two countries' population-weighted relief profiles more similar than two arbitrarily chosen countries in the pool: Slovenia's low-relief valleys and Georgia's Caucasus highlands sit in the same "Mediterranean/Caucasus" macro-grouping used elsewhere in this program only by administrative and cultural convention, not by terrain. This null result, honestly reported exactly as computed, is itself informative: it shows that terrain relief is a genuinely orthogonal source of national heterogeneity, not a proxy for the same regional structure that toponymy already captures.

14.4 Cross-framework correlations

Table 14.2 reports Pearson, Spearman and Kendall correlations between $\mathrm{JSD}_{\mathrm{elev}}$ and each of the three previously established pairwise measures, on the identical 210 country pairs used throughout this program.

Table 14.2: Correlations between $\mathrm{JSD}_{\mathrm{elev}}$ and the three previously established pairwise measures, $n = 210$ pairs.

Comparison	Pearson r (p)	Spearman ρ (p)	Kendall τ (p)
$\mathrm{JSD}_{\mathrm{elev}}$ vs. $\mathrm{JSD}_{\mathrm{topo}}$	0.055 (0.431)	0.113 (0.102)	0.072 (0.119)
$\mathrm{JSD}_{\mathrm{elev}}$ vs. $\mathrm{JSD}_{\mathrm{geo}}$	0.173 (0.012)	0.160 (0.020)	0.108 (0.020)
$\mathrm{JSD}_{\mathrm{elev}}$ vs. shape distance (Chapter 13)	-0.197 (0.004)	-0.166 (0.016)	-0.107 (0.021)

Against $\mathrm{JSD}_{\mathrm{topo}}$, the association is weak and not statistically significant by any of the three measures, echoing the toponymic independence already found for $\mathrm{JSD}_{\mathrm{geo}}$ and for shape distance in the earlier chapters: naming-convention divergence continues to carry no detectable relationship to any of the three genuinely geometric or physical descriptors examined across this four-chapter program. Against $\mathrm{JSD}_{\mathrm{geo}}$ and against shape distance, however, $\mathrm{JSD}_{\mathrm{elev}}$ shows

small but nominally significant associations at the conventional $\alpha = 0.05$ threshold ($p = 0.012$ and $p = 0.004$ respectively by the Pearson test): countries whose internal geodesic-distance distributions are more different from one another also tend to have somewhat more different elevation profiles, and countries whose *shapes* are more different tend, if anything, to have slightly *less* different elevation profiles (a negative association). Neither correlation is large: r^2 is under 4% in both cases, meaning terrain relief divergence explains at most a small fraction of the variance in either measure, and Section 14.4 shows below that neither survives a proper multiple-comparisons correction. As a further single-country-level check (Table 14.1's CV column against E and δ directly, $n = 21$ rather than $n = 210$ pairs), no relationship at all is detectable: elevation coefficient of variation correlates with elongation at $r = 0.139$ ($p = 0.548$) and with isoperimetric deficit at $r = 0.030$ ($p = 0.899$), both entirely consistent with independence.

Multiple Comparisons

Six pairwise correlations across four measures are reported in this chapter and the next (JSD_{topo}, JSD_{geo}, shape distance, JSD_{elev}, taken two at a time), of which two nominally clear the conventional $p < 0.05$ threshold (JSD_{elev} vs. JSD_{geo} and vs. shape distance). A Bonferroni correction for six simultaneous comparisons requires $p < 0.05/6 \approx 0.0083$ for significance; neither correlation clears this stricter bar ($p = 0.012$ and $p = 0.004$ by Pearson, but $p = 0.020$ and $p = 0.016$ by the arguably more appropriate rank-based Spearman test, and $p = 0.020/p = 0.021$ by Kendall's τ, none below 0.0083). The honest reading is therefore that this chapter finds, at most, weak and statistically fragile evidence of association between terrain relief divergence and the two other geometric measures, not a robust cross-framework relationship of the kind that a single, uncorrected $p < 0.05$ might suggest at first glance.

Chapter 15

A Four-Axis Correlation Synthesis and an Exploratory Country Typology

Four genuinely independent pairwise measures of national heterogeneity have now been constructed across Chapters 12 to 14 from the same underlying GeoNames sample of 21 countries: toponymic divergence (JSD_{topo}), internal geodesic-distance-distribution divergence (JSD_{geo}), country-shape distance built from elongation and isoperimetric deficit, and terrain-relief divergence (JSD_{elev}). This closing chapter assembles the complete pairwise correlation structure among all four measures in one place, and uses the four underlying single-country descriptors — E, δ, elevation CV, and median elevation — to ask a fifth, exploratory question left untouched by any of the pairwise tests: do the 21 countries fall into any natural typology when all four descriptors are considered jointly?

Methodological Note

Every correlation reported in this chapter has already been reported, individually, in Chapters 12 to 14; nothing here is a new computation, only a new arrangement of already-executed results into a single matrix for direct comparison, plus one genuinely new computation — the hierarchical clustering of Section 15.2 — run once, on the real standardized descriptor values, with no post-hoc adjustment of the cluster count or the descriptor set.

15.1 The complete correlation matrix

Table 15.1 reports Kendall's τ (chosen over Pearson or Spearman as the primary summary, following the rank-correlation preference already established in Section 9.3) for all six pairs among the four measures, each computed on the same 210 country pairs.

Table 15.1: Kendall's τ rank correlation among all four pairwise measures, $n = 210$ country pairs. Bonferroni-adjusted significance threshold for six simultaneous comparisons: $p < 0.0083$.

Measure 1	Measure 2	Kendall τ	p
JSD_{topo}	JSD_{geo}	-0.008	0.865
JSD_{topo}	shape distance	-0.028	0.550
JSD_{topo}	JSD_{elev}	0.072	0.119
JSD_{geo}	shape distance	-0.103	0.027
JSD_{geo}	JSD_{elev}	0.108	0.020
shape distance	JSD_{elev}	-0.107	0.021

Not one of the six pairwise correlations in Table 15.1 survives the Bonferroni-adjusted threshold of $p < 0.0083$. The overall picture, reading across all four chapters of this cross-framework program together, is therefore one of substantial — though not absolute — *independence*: to-

ponymic divergence shares no detectable relationship with any of the three geometric or physical measures ($|\tau| \leq 0.072$ in all three comparisons against $\mathrm{JSD}_{\mathrm{topo}}$), while the three non-toponymic measures show small, same-order-of-magnitude, and individually non-robust pairwise associations with one another ($0.10 \leq |\tau| \leq 0.11$). No single pair of measures stands out as capturing the same underlying structure by a different route; each of the four appears to isolate a genuinely distinct axis of variation among the 21 countries in the pool. This is, on balance, the strongest and most consistent finding of the entire four-chapter program: four methodologically unrelated ways of measuring how different two countries are from one another — by the names they give places, by the internal shape of their distance structure, by their territorial outline, and by the terrain their populated places sit on — yield four largely uncorrelated rankings, not four dialects of a single underlying "national distinctiveness" scale.

15.2 An exploratory four-descriptor country typology

The pairwise measures of Section 15.1 compare *country pairs*; a complementary, single-country-level question is whether the four underlying descriptors — E, δ, elevation CV, and median elevation, each standardized to zero mean and unit variance across the 21-country pool — group the countries into any recognizable typology. A standard UPGMA (average-linkage) hierarchical clustering of the 21 standardized four-dimensional descriptor vectors, cut at four clusters, gives the assignment in Table 15.2.

Table 15.2: UPGMA clustering (four clusters, average linkage, Euclidean distance) on the standardized four-descriptor vector $(\tilde{E}, \tilde{\delta}, \widetilde{\mathrm{CV}}_{\mathrm{elev}}, \widetilde{\mathrm{median}}_{\mathrm{elev}})$.

Cluster	Countries
1	Morocco, Georgia
2	Slovenia, Turkey, Austria, Switzerland
3	Italy, Spain, Portugal, Croatia, France, Algeria, Tunisia, Libya, Mauritania, Azerbaijan, India, Pakistan, United Kingdom, Ireland
4	Armenia

The honest reading of Table 15.2 is that this four-descriptor space does *not* partition the 21 countries into four balanced, interpretable groups: fourteen of the twenty-one countries — the large majority — fall into a single undifferentiated cluster 3, while clusters 1, 2 and 4 are driven by a small number of countries with extreme values on one or two descriptors rather than by any broad, shared typology. Armenia forms its own singleton cluster 4 essentially because of its outlying elevation profile alone (Table 14.1: mean elevation nearly three times the next-highest country); Morocco and Georgia are grouped in cluster 1 by, respectively, the highest isoperimetric deficit in the pool and one of the higher combinations of elongation and elevation dispersion; and cluster 2's four countries (Slovenia, Turkey, Austria, Switzerland) share moderate-to-high elongation together with comparatively low elevation coefficient of variation, the one grouping in Table 15.2 with a plausible common description — geographically and culturally disparate countries (an Alpine republic, a Balkan state, and a transcontinental one) whose populated-place elevation profiles happen to be comparably *even*, whatever the mean elevation itself. This is precisely the kind of result that should not be over-interpreted: with only 21 data points in four dimensions and no external validation set, a four-cluster split is a description of this particular sample's descriptor geometry, not evidence of any deeper geographic or cultural law, and it is reported here in that spirit, consistent with Section 15.1's finding that the four underlying measures are themselves largely uncorrelated with one another.

15.3 Closing remark on this four-chapter program

Read together, Chapters 12 to 14 and this synthesis chapter constitute a single sustained empirical exercise in testing whether four independently motivated ways of measuring difference among the same 21 countries — naming conventions, internal distance structure, territorial shape, and terrain relief — turn out to be redundant or complementary. The answer, arrived at honestly across four real, executed computations rather than assumed at the outset, is that they are substantially complementary: only weak, statistically fragile associations connect any pair of the four measures, none surviving correction for the six simultaneous comparisons involved, and the one attempt to find a joint typology in Section 15.2 produces a lopsided, largely uninformative partition rather than a clean four-way split. This is not a negative result in any pejorative sense — it is the expected, and arguably the more useful, outcome of deploying several genuinely different quantitative lenses on the same set of objects: each lens keeps showing something the others do not.

Chapter 16

Synthesis and Outlook

16.1 What each chapter actually establishes

It is worth stating plainly, in one place, exactly what mathematical content has been proved or computed in each chapter, since the seven results differ considerably in character.

- **Chapter 2** proves an absolute, unconditional impossibility: no local isometry exists between any curved region of the sphere and the plane (Theorem 2.3), and no projection can be simultaneously conformal and equiareal on a curved region (Proposition 2.4); Tissot's indicatrix (Section 2.5) then supplies the exact closed-form pointwise bookkeeping — scale factors h, k, angular distortion ω, and areal scale s — by which the size of this unavoidable distortion is actually measured at any given point of any given projection. This is the strongest kind of mathematical statement in the monograph — a theorem, not subject to future refutation by cleverer engineering.

- **Chapter 3** proves a sharp optimality characterization within the space of conformal projections: the projection minimizing worst-case scale distortion has constant scale factor on the region's boundary (Theorem 3.1), reducing "best projection for this specific territory" to a solvable boundary-value problem whose existence and uniqueness follow from the Dirichlet principle (Proposition 3.4), and whose distortion is bounded below in a way that grows with the size of the mapped region, per Milnor's 1969 quantitative refinement (Section 3.5).

- **Chapter 4** proves a topological, not metric, obstruction: every continuous tangent vector field on the sphere has at least one zero (Theorem 4.1), with the direct meteorological corollary that no continuous global wind field model can be everywhere nonzero (Corollary 4.4), the further subtlety that this mandatory zero can be projected to the point at infinity and so disappear from a finite polar map (Proposition 4.6), and two fully worked index computations (Section 4.5) showing exactly how the mandatory total index 2 can be realized either as two source/sink-type zeros or as two vortex-type zeros at the poles.

- **Chapter 5** proves a discrete, purely combinatorial analogue of the Gauss–Bonnet theorem: the total angular defect of any genus-zero convex polyhedron is exactly 4π (Theorem 5.2), governing the curvature "budget" available to any polyhedral map projection such as Fuller's Dymaxion map; the exact vertex/edge/face counts of geodesic subdivision (Proposition 5.10) show precisely how that fixed budget is spread over an increasing number of vertices as the reference polyhedron is refined; and Girard's theorem (Theorem 5.4) proves the exact continuous counterpart, showing the same 4π budget reappears as spherical excess summed over the faces of a curved great-circle triangulation, verified explicitly on the icosahedron's own combinatorics in Example 5.7.

- **Chapter 6** proves an explicit closed-form result via ordinary calculus: constant-bearing curves on the sphere are logarithmic spirals about each pole (Proposition 6.3) that wind around infinitely often yet have finite total arc length (Theorem 6.4), are exactly the curves Mercator's projection was constructed to render as straight lines, and admit a closed-form

distance formula between any two latitudes (Corollary 6.5), worked out numerically for an explicit bearing and latitude pair in Example 6.6 and compared quantitatively against the great-circle distance for the same endpoints in Example 6.8.

- **Chapter 7** proves two classical theorems from projective and inversive geometry: the gnomonic projection sends every great circle to a straight line (Theorem 7.2), which is exactly why it underlies the polyhedral projections of Chapter 5, while the stereographic projection is simultaneously conformal (Theorem 7.6) and sends every circle to a circle or line (Theorem 7.7), which is exactly why it renders the loxodromes of Chapter 6 as clean logarithmic spirals about straight radial meridian-images.

- **Chapter 8** presents an original proposal together with an executed, non-invented computation: a Jensen–Shannon-divergence-based distributional distortion index (Definition 8.1), computed on four systematically constructed coordinate grids for twelve standard projections (including Winkel Tripel, Mollweide — whose auxiliary angle is obtained by genuine Newton–Raphson iteration rather than any closed-form shortcut — Aitoff, and, added in this latest expansion, Albers Equal-Area Conic, Hammer, and Van der Grinten, each implemented independently and cross-checked, Remark 8.6), yielding a reproducible ranking that is broadly, but not uniformly, consistent with classical cartographic assessment: equal-area projections are generally favoured and generally stable, Equirectangular and Robinson are generally disfavoured, but the global-grid winner turns out to be Aitoff — neither conformal nor equal-area — which remains the single worst performer at mid-latitude, the sharpest region-dependent reversal recorded in the chapter. Hammer, Aitoff's equal-area sibling built from an otherwise identical azimuthal-halving construction, finishes a close second globally while never collapsing regionally the way Aitoff does, isolating equal-area status as the likely source of that stability; Van der Grinten emerges as a genuinely unspecialized compromise that beats its historical successor Robinson on every sampled grid. This sits alongside a secondary diagnostic — the median-matching scale ratio of Section 8.5 — that independently and quantitatively confirms the textbook scale-inflation behaviour of the Mercator projection toward the poles, and shows Van der Grinten's own scale ratio to be the least stable of the three newly added projections, consistent with its well-known polar exaggeration.

- **Chapter 9** does not introduce a new mathematical result but subjects the claims of Chapter 8 to direct numerical stress-testing, on code built on top of, and executed against, the same pipeline, re-run for all twelve projections: it shows the global-grid winner is Aitoff at three of four tested histogram bin counts but that Hammer edges narrowly ahead of Aitoff at $n_{\text{bins}} = 120$ (Section 9.1) — a genuine, honestly reported instability at the very top of the ranking that did not exist with the previous nine-projection pipeline, where Aitoff won at every bin count tested. The resampling stability check shows Aitoff's global win reproduced in 76.5% of 200 no-replacement subsamples (down from 96% with nine projections, since Hammer now wins a further 20.5% of subsamples) and its mid-latitude weakness remains the modal, though not unanimous, outcome under subsampling (Section 9.2, including a documented exploratory bootstrap variant that was tried first and found to inject its own resampling artifact, Remark 9.1). The Kendall's τ analysis, re-run with all twelve projections, now finds *two* of the six pairwise rank correlations among the four sampling regions statistically significant at the conventional threshold — global-versus-equatorial ($\tau = 0.446$, $p = 0.046$) and mid-latitude-versus-high-latitude ($\tau = 0.485$, $p = 0.031$) — a genuine change from the nine-projection edition, where none reached significance (Section 9.3).

- **Chapter 10** extends the pipeline once more, from twelve to seventeen projections and from four to five sampling grids (Section 10.1), and adds four further, genuinely new diagnostics

rather than retrofitting the old ones. Grid resolution — a robustness axis distinct from bin count — proves unstable at a coarse $20°$ step even though the $5°$ and $10°$ resolutions used throughout the rest of the monograph agree with each other (Section 10.2). A first direct correlation between the IDI and the classical, pointwise Tissot indicatrix finds the association weak and not statistically significant at $n = 17$ (Section 10.3), a substantive negative result consistent with the IDI capturing distributional rather than pointwise structure. A formal subsampling-based significance test of four close pairwise margins, including the Arctic Polyconic/Bonne gap that looked decisive at face value, finds none of them statistically distinguishable from resampling noise (Section 10.4). Alongside these, the new Arctic grid confirms the classical, pointwise expectation that pseudoconical and conic constructions (Polyconic, Bonne, Lambert Conformal Conic) dominate near the pole while Mercator is by far the worst performer there, and Gall–Peters's poor global ranking despite its equal-area property is flagged as a genuine tension left unresolved by the equal-area explanation offered in Section 8.4.

- **Chapter 11** documents, without retyping the implementation line by line, exactly how every numerical claim in Chapters 8 to 10 was produced and checked: the deterministic grid constructions, the haversine geodesic formula validated against two independent known city-pair distances, a three-route verification discipline applied to all seventeen projection formulas (published worked examples, internal algebraic cross-checks, and known closed-form pole values), the Newton–Raphson solver behind Mollweide's auxiliary angle (converging in at most 14 iterations across 500 tested latitudes), and the binning, regularization, and median-rescaling choices underlying the Jensen–Shannon-divergence-based IDI itself.

- **Chapter 12** finally executes the cross-framework test raised as an open question in Section 8.7: using real, per-record GeoNames data for the same twenty-one countries examined empirically, from a different angle, in the companion volume, it finds no detectable correlation between toponymic bigram-divergence and geodesic-distance-structure divergence across 210 country pairs (Pearson $r = 0.011$, p-values above 0.85 throughout), a null result confirmed stable across a four-fold change in geometric sample size and unaffected by substituting equal-area-projected for geodesic distances, even though the toponymic measure alone carries a strong, highly significant regional signal ($p < 0.0002$) that the geometric measure does not share. A region-by-region breakdown further shows *why* the geometric measure misses this regional signal: national-scale heterogeneity within culturally coherent regions (Slovenia beside France, Armenia beside Turkey) is large enough to swamp it. The two JSD-based frameworks are shown, empirically rather than merely by stated analogy, to be complementary rather than redundant.

- **Chapter 14** adds a fourth pairwise measure, $\mathrm{JSD}_{\mathrm{elev}}$, built from the real DEM elevation field attached to the same GeoNames records and reusing the identical 350-place sample throughout, finding no regional signal at all (same-region versus cross-region difference consistent with chance, $p = 0.328$) and only weak associations with $\mathrm{JSD}_{\mathrm{geo}}$ and with the shape distance of Chapter 13 ($|r| < 0.20$ in both cases) that do not survive a Bonferroni correction for the six comparisons made across the two chapters, while remaining, like every other measure examined, undetectably related to toponymic divergence.

- **Chapter 15** assembles the complete Kendall-τ correlation matrix among all four measures constructed across the preceding three chapters, finding none of the six pairwise correlations significant after correction for multiple comparisons, and reports one further, genuinely new and explicitly exploratory computation: a four-descriptor hierarchical clustering of the twenty-one countries that produces a lopsided partition (fourteen of twenty-one countries in a single undifferentiated cluster) rather than a clean typology, honestly pre-

sented as a description of this particular sample rather than as evidence of any deeper geographic law.

16.2 Two recurring invariants: $\chi(S^2) = 2$ and the impossibility gap

Two threads recur across otherwise disparate chapters. First, the Euler characteristic of the sphere, $\chi(S^2) = 2$, appears three times in unrelated-looking guises: as the index sum forced on any vector field by the Poincaré–Hopf theorem (Chapter 4); as $4\pi = 2\pi\chi$, the total angular defect forced on any polyhedron by Descartes's theorem (Chapter 5); and again as $4\pi = 2\pi\chi$, the total spherical excess forced on any great-circle triangulation of the sphere itself by Girard's theorem (Section 5.3). These are not coincidentally similar facts but three applications of the same underlying identity, $\int_M K\, dA = 2\pi\chi(M)$: one purely topological and continuous (vector field indices), one discrete and combinatorial (polyhedral vertex defects), and one continuous and metric (spherical triangle excess).

Second, the impossibility established in Chapter 2 is not merely a negative result standing apart from the rest of the monograph; it is the shared premise that makes the remaining chapters interesting questions rather than solved problems. Chebyshev's theorem (Chapter 3) is only a nontrivial optimization because perfect isometry is unavailable. The entire family of polyhedral projections in Chapter 5 exists only because no single flat unfolding of the whole sphere is isometric, so the sphere must be cut into pieces, each still individually imperfect, and the discrete Gauss–Bonnet theorem tells us precisely how much curvature "deficit" those pieces must jointly account for. And the very existence of a meaningful, nonzero Informational Distortion Index in Chapter 8 is a distributional restatement, in information-theoretic language, of the same impossibility: if isometric projection were possible, every projection's IDI relative to the true geodesic distribution would be exactly zero on every region and every sampling grid, and Chapter 8 would have nothing to measure.

16.3 Directions not pursued here

Several natural extensions of this material were consciously left outside the scope of this monograph, and are recorded here for completeness rather than pursued:

- A fully rigorous, general-region (not merely small-region-linearized) treatment of Chebyshev's theorem requires the modern theory of the Beltrami equation and quasiconformal mappings, referenced but not developed in Chapter 3.
- The interaction between the Hairy Ball obstruction (Chapter 4) and higher-order vector field data (e.g., vorticity, not merely the raw velocity field) was not explored, though it is a natural next question for a genuinely atmospheric-science-oriented treatment.
- The Informational Distortion Index of Chapter 8 was computed here only on systematically constructed coordinate grids, precisely to avoid any risk of invented or misremembered empirical data; applying the same pipeline to genuine GeoNames-derived population centres, to test whether the IDI ranking of projections is sensitive to non-uniform, population-weighted sampling rather than a uniform lattice, was exactly the well-defined follow-up computation identified here — and was subsequently carried out in full as the companion volume, *Empirical Informational Distortion of Map Projections*, rather than folded into this one.

- The joint toponymic-cartographic JSD analysis sketched at the end of Chapter 8 (Section 8.7) is no longer an open direction: Chapter 12 carries it out with real GeoNames data and finds the two frameworks statistically independent, a result summarized above.
- The one follow-up Section 12.7 left unattempted — a geometric measure built directly from country-shape descriptors (compactness, elongation) rather than from internal pairwise-distance distributions — is no longer open: Chapter 13 constructs exactly such a pair of descriptors and repeats the cross-framework test against them, finding the same independence from toponymic structure and a weak, explicitly flagged association with the distance-distribution-based measure that does not survive multiple-comparison correction.
- A natural fourth axis — terrain relief, entirely absent from the toponymic, distance-structure and shape measures alike — is likewise no longer open: Chapter 14 builds $\mathrm{JSD}_{\mathrm{elev}}$ directly from real DEM data and Chapter 15 places it alongside the other three measures in a single correlation matrix, closing this four-chapter cross-framework program with a consistent finding of substantial mutual independence among all four measures.

16.4 Closing remark

The seven results gathered here share a common lesson, worth stating without overstatement: cartography's oldest and most familiar complaint — "every map lies a little" — is, on inspection, not one fact but several genuinely different mathematical facts, arising from differential geometry, variational calculus, algebraic topology, discrete combinatorics, classical integral calculus, projective and inversive geometry, and information theory respectively, each with its own precise statement, its own proof, and its own specific cartographic consequence. Treating "map distortion" as a single undifferentiated phenomenon obscures how much rigorous, provable structure actually governs the many different ways in which it is impossible to draw the Earth.

Appendix A

Full Pairwise Shape-Distance and Divergence Table

This appendix lists, for full transparency and independent reproducibility, the complete set of $\binom{21}{2} = 210$ country pairs underlying Chapter 13: the two single-descriptor shape distances $|\Delta E|$ and $|\Delta\delta|$ of Definitions 13.1 and 13.2, the combined standardized shape distance of Section 13.1, and the toponymic and geodesic-structure divergences $\mathrm{JSD}_{\mathrm{topo}}$ and $\mathrm{JSD}_{\mathrm{geo}}$ carried over unchanged from Chapter 12. All values are the direct, executed output of the same pipeline described in Section 13.1 and Section 12.2; nothing in this table is hand-typed or retrofitted. Pairs are ordered alphabetically by the name of country A for ease of lookup.

Table A.1: All 210 pairwise shape-distance and divergence values underlying Sections 13.3 and 13.4: elongation difference $|\Delta E|$, deficit difference $|\Delta\delta|$, combined standardized shape distance, $\mathrm{JSD}_{\mathrm{topo}}$, and $\mathrm{JSD}_{\mathrm{geo}}$ (all divergences in bits).

Country A	Country B	$\lvert\Delta E\rvert$	$\lvert\Delta\delta\rvert$	Comb.	$\mathrm{JSD}_{\mathrm{topo}}$	$\mathrm{JSD}_{\mathrm{geo}}$
Algeria	Armenia	0.140	0.054	0.603	0.3807	0.5233
Algeria	Austria	0.769	0.121	1.763	0.3573	0.2482
Algeria	Azerbaijan	0.104	0.072	0.762	0.3484	0.2701
Algeria	Georgia	2.319	0.197	4.254	0.3474	0.2220
Algeria	India	0.043	0.041	0.429	0.3162	0.2897
Algeria	Ireland	0.072	0.037	0.399	0.3701	0.3682
Algeria	Libya	0.093	0.036	0.399	0.3915	0.0936
Algeria	Mauritania	0.451	0.011	0.736	0.2083	0.0249
Algeria	Morocco	1.280	0.304	3.758	0.0724	0.0355
Algeria	Pakistan	1.324	0.052	2.198	0.2633	0.0271
Algeria	Switzerland	0.272	0.020	0.484	0.3211	0.4755
Algeria	Tunisia	0.399	0.030	0.714	0.0559	0.2457
Algeria	Turkey	1.073	0.096	1.991	0.3651	0.0545
Algeria	United Kingdom	0.584	0.015	0.953	0.3455	0.1501
Armenia	Austria	0.629	0.067	1.228	0.3481	0.2265
Armenia	Azerbaijan	0.037	0.126	1.304	0.2114	0.2614
Armenia	Georgia	2.179	0.143	3.807	0.2296	0.2285
Armenia	India	0.098	0.095	0.995	0.2400	0.8332
Armenia	Ireland	0.212	0.091	1.001	0.3278	0.1334
Armenia	Pakistan	1.183	0.003	1.907	0.2762	0.5768
Armenia	Switzerland	0.132	0.074	0.792	0.3310	0.0261
Armenia	Turkey	0.932	0.042	1.563	0.2670	0.7014
Armenia	United Kingdom	0.443	0.069	1.010	0.3349	0.3268
Austria	Ireland	0.841	0.158	2.124	0.2864	0.0450
Austria	Switzerland	0.497	0.141	1.663	0.0906	0.1462
Austria	United Kingdom	0.186	0.136	1.440	0.2143	0.0536
Azerbaijan	Austria	0.666	0.193	2.267	0.3724	0.0165
Azerbaijan	Georgia	2.215	0.269	4.522	0.2393	0.0238
Azerbaijan	India	0.061	0.031	0.334	0.2177	0.6940
Azerbaijan	Ireland	0.176	0.035	0.459	0.3043	0.0539

Country A	Country B	$\lvert\Delta E\rvert$	$\lvert\Delta\delta\rvert$	Comb.	$\mathrm{JSD}_{\mathrm{topo}}$	$\mathrm{JSD}_{\mathrm{geo}}$
Azerbaijan	Pakistan	1.220	0.124	2.343	0.2496	0.3167
Azerbaijan	Switzerland	0.168	0.052	0.604	0.3233	0.1813
Azerbaijan	Turkey	0.969	0.168	2.332	0.1497	0.4600
Azerbaijan	United Kingdom	0.480	0.057	0.972	0.3303	0.0628
Croatia	Algeria	0.267	0.011	0.443	0.4096	0.2863
Croatia	Armenia	0.407	0.065	0.936	0.3281	0.2508
Croatia	Austria	1.036	0.132	2.154	0.4082	0.0403
Croatia	Azerbaijan	0.370	0.061	0.871	0.3129	0.0139
Croatia	France	0.201	0.018	0.371	0.3622	0.3309
Croatia	Georgia	2.586	0.207	4.684	0.2763	0.0371
Croatia	India	0.309	0.031	0.590	0.3201	0.7038
Croatia	Ireland	0.195	0.026	0.416	0.3631	0.0580
Croatia	Libya	0.174	0.025	0.383	0.5661	0.3991
Croatia	Mauritania	0.717	0.022	1.178	0.3982	0.3217
Croatia	Morocco	1.547	0.315	4.095	0.4398	0.2088
Croatia	Pakistan	1.590	0.062	2.642	0.3511	0.3338
Croatia	Slovenia	1.286	0.074	2.208	0.0669	0.2738
Croatia	Switzerland	0.539	0.009	0.873	0.3464	0.1830
Croatia	Tunisia	0.666	0.040	1.151	0.3736	0.0342
Croatia	Turkey	1.339	0.106	2.421	0.2822	0.4744
Croatia	United Kingdom	0.850	0.005	1.371	0.3795	0.0734
France	Algeria	0.066	0.028	0.309	0.3446	0.0675
France	Armenia	0.206	0.082	0.912	0.3848	0.5726
France	Austria	0.835	0.150	2.048	0.2922	0.2848
France	Azerbaijan	0.170	0.044	0.529	0.3591	0.3136
France	Georgia	2.385	0.225	4.491	0.3701	0.2497
France	India	0.109	0.013	0.220	0.3716	0.4374
France	Ireland	0.006	0.009	0.092	0.2410	0.4224
France	Libya	0.027	0.008	0.091	0.5552	0.1664
France	Mauritania	0.517	0.040	0.928	0.3112	0.0304
France	Morocco	1.346	0.332	4.059	0.3687	0.0435
France	Pakistan	1.390	0.080	2.386	0.3686	0.0813
France	Switzerland	0.338	0.008	0.552	0.1768	0.5272
France	Tunisia	0.465	0.058	0.960	0.3180	0.2828
France	Turkey	1.139	0.124	2.236	0.3493	0.1147
France	United Kingdom	0.650	0.013	1.056	0.2012	0.1517
Georgia	Austria	1.550	0.075	2.616	0.4013	0.0168
Georgia	India	2.276	0.238	4.414	0.2393	0.6616
Georgia	Ireland	2.391	0.234	4.547	0.3784	0.0819
Georgia	Pakistan	0.995	0.145	2.196	0.2782	0.2824
Georgia	Switzerland	2.047	0.217	3.985	0.3436	0.1747
Georgia	Turkey	1.246	0.101	2.263	0.3014	0.4078
Georgia	United Kingdom	1.735	0.212	3.550	0.3870	0.0414
India	Austria	0.727	0.162	2.045	0.3524	0.6735
India	Ireland	0.115	0.004	0.190	0.3159	0.7519
India	Pakistan	1.281	0.093	2.275	0.1388	0.2995
India	Switzerland	0.229	0.021	0.430	0.3277	0.8126
India	United Kingdom	0.541	0.026	0.912	0.3254	0.5919
Italy	Algeria	1.097	0.032	1.798	0.3592	0.0400
Italy	Armenia	0.956	0.086	1.780	0.3820	0.5294
Italy	Austria	0.327	0.154	1.670	0.3651	0.2806
Italy	Azerbaijan	0.993	0.040	1.652	0.3460	0.3091
Italy	Croatia	1.363	0.022	2.208	0.2505	0.3263
Italy	France	1.163	0.004	1.874	0.2248	0.0586
Italy	Georgia	1.222	0.229	3.077	0.2936	0.2669
Italy	India	1.054	0.009	1.701	0.3265	0.3485

Country A	Country B	$\lvert\Delta E\rvert$	$\lvert\Delta\delta\rvert$	Comb.	$\mathrm{JSD}_{\mathrm{topo}}$	$\mathrm{JSD}_{\mathrm{geo}}$
Italy	Ireland	1.169	0.005	1.884	0.2839	0.3845
Italy	Libya	1.189	0.004	1.917	0.5480	0.0627
Italy	Mauritania	0.646	0.044	1.134	0.3767	0.0176
Italy	Morocco	0.183	0.336	3.485	0.3948	0.0671
Italy	Pakistan	0.227	0.084	0.937	0.3314	0.0186
Italy	Portugal	0.848	0.193	2.418	0.1642	0.2728
Italy	Slovenia	0.077	0.095	0.992	0.2686	0.5455
Italy	Spain	0.430	0.234	2.514	0.1185	0.1012
Italy	Switzerland	0.825	0.012	1.335	0.2534	0.4797
Italy	Tunisia	0.698	0.062	1.294	0.3255	0.2808
Italy	Turkey	0.024	0.128	1.320	0.3429	0.0664
Italy	United Kingdom	0.513	0.017	0.845	0.2828	0.1841
Libya	Armenia	0.233	0.090	1.003	0.4040	0.4841
Libya	Austria	0.862	0.157	2.137	0.5217	0.3576
Libya	Azerbaijan	0.196	0.036	0.489	0.3388	0.3895
Libya	Georgia	2.412	0.233	4.569	0.4834	0.3596
Libya	India	0.135	0.005	0.225	0.3619	0.2804
Libya	Ireland	0.021	0.001	0.035	0.4768	0.4043
Libya	Mauritania	0.544	0.047	1.003	0.4514	0.1007
Libya	Pakistan	1.416	0.087	2.454	0.2906	0.0554
Libya	Switzerland	0.365	0.016	0.611	0.5089	0.4508
Libya	Turkey	1.166	0.132	2.318	0.4652	0.0838
Libya	United Kingdom	0.676	0.021	1.111	0.4762	0.2951
Mauritania	Armenia	0.311	0.043	0.667	0.3377	0.5435
Mauritania	Austria	0.318	0.110	1.245	0.2954	0.2678
Mauritania	Azerbaijan	0.347	0.084	1.028	0.2960	0.2974
Mauritania	Georgia	1.868	0.185	3.567	0.3080	0.2443
Mauritania	India	0.408	0.053	0.852	0.2754	0.3646
Mauritania	Ireland	0.523	0.048	0.980	0.3085	0.3951
Mauritania	Pakistan	0.873	0.040	1.466	0.3014	0.0257
Mauritania	Switzerland	0.179	0.031	0.432	0.2557	0.4965
Mauritania	Turkey	0.622	0.084	1.327	0.2986	0.0635
Mauritania	United Kingdom	0.133	0.027	0.348	0.2910	0.1549
Morocco	Armenia	1.140	0.250	3.168	0.3741	0.5077
Morocco	Austria	0.511	0.183	2.059	0.3897	0.1722
Morocco	Azerbaijan	1.176	0.376	4.322	0.3243	0.1867
Morocco	Georgia	1.039	0.107	2.008	0.3677	0.1518
Morocco	India	1.237	0.345	4.084	0.3136	0.4147
Morocco	Ireland	1.352	0.341	4.142	0.3677	0.3022
Morocco	Libya	1.373	0.340	4.150	0.3265	0.1704
Morocco	Mauritania	0.829	0.293	3.305	0.2265	0.0356
Morocco	Pakistan	0.044	0.253	2.610	0.2563	0.0674
Morocco	Switzerland	1.008	0.324	3.718	0.3588	0.4396
Morocco	Tunisia	0.881	0.274	3.168	0.0902	0.1695
Morocco	Turkey	0.207	0.209	2.179	0.3695	0.1276
Morocco	United Kingdom	0.696	0.319	3.482	0.3510	0.0759
Pakistan	Austria	0.554	0.070	1.148	0.3493	0.3015
Pakistan	Ireland	1.396	0.088	2.428	0.3050	0.4137
Pakistan	Switzerland	1.052	0.071	1.847	0.3115	0.5263
Pakistan	United Kingdom	0.740	0.067	1.376	0.2836	0.2030
Portugal	Algeria	0.248	0.161	1.711	0.3507	0.1986
Portugal	Armenia	0.108	0.107	1.118	0.4199	0.2611
Portugal	Austria	0.521	0.040	0.934	0.3761	0.0825
Portugal	Azerbaijan	0.145	0.233	2.418	0.3765	0.0837
Portugal	Croatia	0.515	0.172	1.957	0.3214	0.0706
Portugal	France	0.314	0.189	2.018	0.2461	0.2993

Country A	Country B	$\lvert\Delta E\rvert$	$\lvert\Delta\delta\rvert$	Comb.	$\mathrm{JSD}_{\mathrm{topo}}$	$\mathrm{JSD}_{\mathrm{geo}}$
Portugal	Georgia	2.071	0.036	3.357	0.3550	0.0688
Portugal	India	0.206	0.202	2.113	0.3628	0.5158
Portugal	Ireland	0.320	0.198	2.109	0.3234	0.1222
Portugal	Libya	0.341	0.197	2.107	0.5418	0.3068
Portugal	Mauritania	0.202	0.150	1.579	0.3457	0.2667
Portugal	Morocco	1.032	0.143	2.224	0.3847	0.1772
Portugal	Pakistan	1.075	0.110	2.069	0.3547	0.2691
Portugal	Slovenia	0.771	0.098	1.601	0.3188	0.2828
Portugal	Switzerland	0.024	0.181	1.868	0.2954	0.2088
Portugal	Tunisia	0.151	0.131	1.376	0.3212	0.0816
Portugal	Turkey	0.824	0.065	1.490	0.3857	0.3603
Portugal	United Kingdom	0.335	0.176	1.897	0.3254	0.0983
Slovenia	Algeria	1.019	0.063	1.768	0.4038	0.5421
Slovenia	Armenia	0.879	0.009	1.420	0.3460	0.0054
Slovenia	Austria	0.250	0.058	0.723	0.3949	0.2384
Slovenia	Azerbaijan	0.916	0.135	2.032	0.3559	0.2780
Slovenia	France	1.085	0.091	1.987	0.3588	0.5919
Slovenia	Georgia	1.300	0.134	2.508	0.3258	0.2464
Slovenia	India	0.977	0.104	1.907	0.3566	0.8472
Slovenia	Ireland	1.091	0.100	2.041	0.3679	0.1442
Slovenia	Libya	1.112	0.099	2.064	0.5736	0.4965
Slovenia	Mauritania	0.569	0.052	1.061	0.4054	0.5598
Slovenia	Morocco	0.261	0.241	2.523	0.4372	0.5260
Slovenia	Pakistan	0.304	0.012	0.505	0.3624	0.5953
Slovenia	Switzerland	0.747	0.083	1.478	0.3392	0.0283
Slovenia	Tunisia	0.620	0.033	1.057	0.3707	0.2360
Slovenia	Turkey	0.053	0.032	0.346	0.3225	0.7207
Slovenia	United Kingdom	0.436	0.078	1.071	0.3671	0.3460
Spain	Algeria	0.667	0.202	2.346	0.3161	0.0804
Spain	Armenia	0.527	0.148	1.747	0.3613	0.6430
Spain	Austria	0.102	0.081	0.849	0.3378	0.3769
Spain	Azerbaijan	0.563	0.274	2.971	0.3163	0.4049
Spain	Croatia	0.934	0.213	2.661	0.2820	0.4124
Spain	France	0.733	0.230	2.654	0.1851	0.0618
Spain	Georgia	1.652	0.005	2.663	0.3220	0.3354
Spain	India	0.624	0.243	2.704	0.3189	0.2842
Spain	Ireland	0.739	0.239	2.740	0.2568	0.5140
Spain	Libya	0.760	0.238	2.744	0.4816	0.1776
Spain	Mauritania	0.216	0.191	1.998	0.3196	0.0780
Spain	Morocco	0.613	0.102	1.445	0.3460	0.0943
Spain	Pakistan	0.657	0.151	1.880	0.3091	0.1092
Spain	Portugal	0.419	0.041	0.796	0.0858	0.3257
Spain	Slovenia	0.352	0.139	1.541	0.2825	0.6641
Spain	Switzerland	0.395	0.222	2.376	0.2324	0.6097
Spain	Tunisia	0.268	0.172	1.828	0.2788	0.3747
Spain	Turkey	0.406	0.106	1.278	0.3119	0.1146
Spain	United Kingdom	0.083	0.217	2.245	0.2641	0.2373
Switzerland	Ireland	0.344	0.017	0.582	0.2236	0.0598
Switzerland	United Kingdom	0.312	0.005	0.504	0.1561	0.2624
Tunisia	Armenia	0.259	0.024	0.486	0.3697	0.2230
Tunisia	Austria	0.370	0.091	1.117	0.3419	0.0039
Tunisia	Azerbaijan	0.295	0.102	1.155	0.3293	0.0114
Tunisia	Georgia	1.920	0.167	3.541	0.3443	0.0118
Tunisia	India	0.356	0.071	0.931	0.3025	0.6731
Tunisia	Ireland	0.471	0.067	1.026	0.3429	0.0496
Tunisia	Libya	0.492	0.066	1.044	0.3729	0.3582

Country A	Country B	$\lvert\Delta E\rvert$	$\lvert\Delta\delta\rvert$	Comb.	$\mathrm{JSD}_{\mathrm{topo}}$	$\mathrm{JSD}_{\mathrm{geo}}$
Tunisia	Mauritania	0.052	0.018	0.208	0.2083	0.2668
Tunisia	Pakistan	0.925	0.022	1.507	0.2379	0.2980
Tunisia	Switzerland	0.127	0.050	0.552	0.2966	0.1498
Tunisia	Turkey	0.674	0.066	1.280	0.3453	0.4355
Tunisia	United Kingdom	0.185	0.045	0.552	0.3137	0.0496
Turkey	Austria	0.304	0.026	0.556	0.3565	0.4385
Turkey	India	1.030	0.137	2.179	0.2746	0.2236
Turkey	Ireland	1.145	0.133	2.298	0.3166	0.5666
Turkey	Pakistan	0.251	0.044	0.610	0.3396	0.0341
Turkey	Switzerland	0.801	0.115	1.756	0.2989	0.6667
Turkey	United Kingdom	0.489	0.111	1.389	0.3391	0.3063
United Kingdom	Ireland	0.656	0.022	1.081	0.1036	0.1428

References

The classical theorems and historical projections discussed throughout this monograph draw on the following primary and standard secondary sources. Biographical and historical attributions for the individual named projections (invention dates, inventors, adoption history) not otherwise footnoted follow the historical accounts in [Snyder, 1993] and [Snyder & Voxland, 1989] throughout.

Bibliography

[Bonnet, 1848] Bonnet, P. O. (1848). Mémoire sur la théorie générale des surfaces. *Journal de l'École Polytechnique*, Tome 19 (Cahier 32), 1–146.

[Chebyshev, 1856] Chebyshev, P. L. (1856). Sur la construction des cartes géographiques. *Bulletin de la Classe Physico-Mathématique de l'Académie Impériale des Sciences de Saint-Pétersbourg*, Tome 14, 257–261. Reprinted in *Œuvres de P. L. Tchebychef*, Vol. 1, ed. A. Markoff and N. Sonin (St. Petersburg, 1899), 233–236.

[Fuller, 1943] Fuller, R. B. (1943). Dymaxion world. *Life Magazine*, March 1, 1943.

[Gauss, 1828] Gauss, C. F. (1828). Disquisitiones generales circa superficies curvas. *Commentationes Societatis Regiae Scientiarum Gottingensis Recentiores*, Vol. 6, 99–146. Presented to the Royal Society of Göttingen on 8 October 1827.

[Girard, 1629] Girard, A. (1629). *Invention nouvelle en l'algèbre*. Amsterdam: Guillaume Iansson Blaeuw.

[Milnor, 1969] Milnor, J. (1969). A problem in cartography. *American Mathematical Monthly*, 76(10), 1101–1112.

[Snyder, 1987] Snyder, J. P. (1987). *Map Projections: A Working Manual*. U.S. Geological Survey Professional Paper 1395. Washington, DC: U.S. Government Printing Office.

[Snyder, 1993] Snyder, J. P. (1993). *Flattening the Earth: Two Thousand Years of Map Projections*. Chicago: University of Chicago Press.

[Snyder & Voxland, 1989] Snyder, J. P., & Voxland, P. M. (1989). *An Album of Map Projections*. U.S. Geological Survey Professional Paper 1453. Washington, DC: U.S. Government Printing Office.